ETH
Swiss Federal Institute of Technology Zurich

Seminar for Statistics

Department of Mathematics

Semester Project Autumn 2015

arXiv:1704.00575v1 [stat.AP] 3 Apr 2017

Emiliano Díaz

Sparse Mean Localization by Information Theory

Submission Date: April 26th 2016

Adviser Nico Gorbach, Stefan Bauer
Supervisor: Prof. Dr. Markus Kalisch, Prof. Joachim M. Buhmann

I would like to express my gratitude to Stefan Bauer and Nico Gorbach who gave me all the support I needed, to Prof. Dr. Joachim Buhmann for the opportunity to work in such as a fascinating topic, and to Dr. Markus Kalisch for all the support, in this semester project and throughout the Masters' programme.

Abstract

Sparse feature selection is necessary when we fit statistical models, we have access to a large group of features, don't know which are relevant, but assume that most are not. Alternatively, when the number of features is larger than the available data the model becomes overparametrized and the sparse feature selection task involves selecting the most informative variables for the model. When the model is a simple location model and the number of relevant features does not grow with the total number of features, sparse feature selection corresponds to sparse mean estimation. We deal with a simplified mean estimation problem consisting of an additive model with gaussian noise and mean that is in a restricted, finite hypothesis space (parameter space). This restriction simplifies the mean estimation problem into a selection problem of combinatorial nature. Although the hypothesis space is finite, its size is exponential in the dimension of the mean. In limited data settings and when the size of the hypothesis space depends on the amount of data or on the dimension of the data, choosing an approximation set of hypotheses is a desirable approach. Choosing a set of hypotheses instead of a single one implies replacing the bias-variance trade off with a resolution-stability trade off. Generalization capacity provides a resolution selection criterion based on allowing the learning algorithm to communicate the largest amount of information in the data to the learner without error. In this work the theory of approximation set coding and generalization capacity is explored in order to understand this approach. We then apply the generalization capacity criterion to the simplified sparse mean estimation problem and detail an importance sampling algorithm which at once solves the difficulty posed by large hypothesis spaces and the slow convergence of uniform sampling algorithms (caused by the skewed distribution of hypothesis costs). Finally we explore how the generalization capacity criterion can be a applied to a more realistic version of the sparse feature selection problem where the number of relevant features grows with the total number of features.

Contents

1 Introduction **1**
1.1 Structure . . . 1
1.2 Problem statement . . . 2
1.3 Pattern analysis . . . 3
1.4 Approximation sets . . . 5
1.5 Generalization capacity . . . 7
1.6 GC and Shannon's noisy-channel coding theorem . . . 10

2 Mean localization **27**
2.1 Generalization capacity . . . 27
2.2 Exhaustive algorithm . . . 29
2.3 Simulation results . . . 29
2.4 Log-sum-exp trick . . . 32
2.5 Variance reduction with common random numbers . . . 33

3 Sparse mean localization **37**
3.1 Exhaustive algorithm . . . 37
3.2 Simulation results: exhaustive algorithm . . . 38
3.3 Sampling algorithm . . . 40
3.4 Simulation results: sampling algorithm . . . 41
3.5 Importance sampling algorithm . . . 44
3.6 Simulation results: importance sampling algorithm . . . 48

4 Towards sparse feature selection for correlated features **53**

5 Summary **59**
5.1 Future work . . . 60

Bibliography **61**

Appendices **61**

Appendix A R Code **63**
A.1 Functions . . . 63
A.2 Script . . . 78

List of Figures

1.1 Boltzmann smoothing: smoothing of costs with Boltzmann weight function 8
1.2 Discrete channel: schematic of input and output of a discrete channel . . . 12
1.3 Communication channel: schematic of communication channel 13
1.4 Shannon code: creation of codebook 14
1.5 Shannon code: decoding 15
1.6 Shannon code: towards capacity 15
1.7 Idealized learning protocol 16
1.8 Learning protocol 20
1.9 Learning protocol: creation of codebook 21
1.10 Learning protocol: decoding with high resolution 21
1.11 Learning protocol: decoding with low resolution 22
1.12 Learning protocol: maximizing learning rate 22

2.1 Component-wise Gibbs distribution 30
2.2 Average information content 31
2.3 Generalization capacity 32
2.4 Calculation of generalization capacity with CRNs 36

3.1 Average information content for different k 38
3.2 Average information content for different σ 39
3.3 Generalization capacity for different k 39
3.4 Generalization capacity estimation for different r 42
3.5 Generalization capacity estimation for different d and r 43
3.6 Importance and frequency of hypotheses 45
3.7 % of hypotheses $\mu \in \mathcal{C}^k$ such that $h > 0$ for $k = 4$ 46
3.8 Generalization capacity estimation for different r 48
3.9 Generalization capacity estimation for different d and r 49
3.10 Generalization capacity estimation for different d and σ 51

Chapter 1

Introduction

It is often the case that when fitting statistical models, the majority of available features are not informative in the sense of the underlying learning task. In other cases the limited amount of data available implies that most features can't be used, even if they are all informative, because the model becomes overparametrized. In both instances sparse feature selection must be done prior or simultaneous to model fitting. In this work we deal with the sparse feature selection problem as it applies to a simplified location model. We first assume the number of relevant features is small and fixed and then explore the case where the number of relevant features is small but grows with the total number of features. Although this problem is well known and studied, for example in **?**, we are interested in how we can apply approximation set coding and generalization capacity to localize the hypothesis class to an *optimal* resolution.

1.1 Structure

The report is organized as follows. Section 1.2 gives a description of the problem we will focus on: sparse mean estimation and sparse feature selection. We want to solve this problem using the *approximation set coding* and *generalization capacity* methodology proposed by **?**, so in Sections 1.3-1.6 we give an introduction to the theory involved. Section 1.3 introduces the *pattern analysis* framework for learning problems. In section 1.4 we explore how, by defining *approximation sets* of hypotheses instead of proposing a single hypothesis as the solution, we are able to move from the normal bias-variance trade-off of learning problems to a resolution-stability trade-off. In section 1.5, with the help of concepts from Sections 1.3 and 1.4, we define various information theoretic concepts such as Boltzmann weights, Gibbs distributions and partition functions, culminating in the definition of *generalization capacity*. We try to give an intuitive understanding of each of these concepts except that of generalization capacity itself. In Section 1.6 we motivate the concept of generalization capacity in analogy to Shannon's noisy channel coding theorem from which it is derived.

In Chapter 2 we estimate the generalization capacity of the squared loss based, empirical risk function for the non-sparse version of the mean localization problem. We concentrate on low-dimensional cases. In Section 2.1 the information theoretic concepts defined in Section 1.5 are applied to the problem at hand culminating in an expression for the

generalization capacity that suggests an *exhaustive* simulating algorithm for its estimation. Section 2.2 includes the pseudo code for implementing this algorithm. Section 2.3 includes the results of implementing the exhaustive simulating algorithm to estimating generalization capacity. Section 2.4 is a note on how to avoid *underflow* problems when implementing this algorithm. In Section 2.5 we explore different ways in which we may incoprorate the use of *common random numbers* into our algorithm as a variance reduction technique.

In Chapter 3 we estimate the generalization capacity for the sparse mean localization problem. Section 3.1 describes the changes and additional tools necessary to implement the algorithm described in 2.2 to the sparse version of the problem. Section 3.2 includes the results of implementing this algorithm to estimating generalization capacity. Since the algorithm will be shown to be inadequate in the high dimensional case, in Section 3.3 we describe a sampling algorithm based on a re-expression of the generalization capacity. Section 3.4 includes the results of this sampling algorithm. This algorithm will be shown to converge too slowly in the number of simulations and so in section 3.5 we describe an *importance* sampling algorithm for estimating generalization capacity. Section 3.6 includes the results of this algorithm.

Chapter 4 is a brief exploration into a more realistic version of sparse feature selection where the number of relevant features grows with the total number of features. We describe the problem and explore some of the difficulties of estimating generalization capacity with a simulation algorithm in this case.

Chapter 5 includes a summary of the report and a list of possible related avenues of future research.

1.2 Problem statement

We deal with the statistical model studied in **?**:

$$X_i = \mu^0 + \epsilon_i$$

where

- $\mu^0 \in \mathbb{B}^d = \{0,1\}^d$
- $X_i, \epsilon_i \in \mathbb{R}^d$
- $\epsilon \sim N(0, \sigma^2 I_d)$
- observations X_i with $i = 1, ..., n$ are i.i.d.

In the general case estimating μ^0 corresponds to selecting a hypothesis μ from the hypothesis space $\mathbb{B}^d$ which has cardinality 2^d. While we first deal with this problem, we will be more interested in a modified version of this problem where:

i.) $||\mu^0||_1 = k$

ii.) $\mu^0 \in \mathbb{B}^d_k = \{\mu \in \mathbb{B}^d : ||\mu||_1 = k\}$

iii.) $|\mathbb{B}^d_k| = \binom{d}{k}$

iv.) It is assumed that k is known.

We first deal with the general case where $k, d \in \mathbb{Z}^+$, $k \leq d$, and then with a *sparse* case where k is kept constant and d grows toward infinity. In Chapter 4 we briefly discuss another *sparsity* condition where $k \approx \frac{d}{\log(d)}$.

1.3 Pattern analysis

Although the classical framework of parameter inference, in which estimators $\hat{\theta}(\cdot)$ are maps from a sample space $\mathcal{X}$ to a parameter space Θ, is appropriate for the problem at hand we introduce *Aproximation Set Coding* (ASC) and *Generalization Capacity* (GC) within the framework of *Pattern Analysis* since they are more relevant in this wider context. The rest of this introductory chapter follows **?** closely.

The problem described in Section 1.2 belongs to the class of problems which are the object of *Pattern Analysis*. The goal of pattern analysis is to map a set of object configurations to a pattern space. Concretely, we want to choose a hypothesis $c \in \mathcal{C}(O^{(n)})$ where:

- $O_i \in \mathcal{O}$ are objects in an object space.
- $O^{(n)} = \{O_1, ..., O_n\} \in \mathcal{O}^{(n)}$ are object sets.
- $c : \mathcal{O}^{(n)} \to \mathcal{P}$ is a hypothesis in a hypothesis class $\mathcal{C}(O^{(n)})$ and $\mathcal{P}$ is a pattern space.

A few remarks about this framework:

i.) The hypothesis class $\mathcal{C}(O^{(n)})$ may or may not depend on the object set. Specifically, the size of the hypothesis class may depend on the object set or not.

ii.) In this exposition the objects in the object set $O_i \in \mathcal{O}$ may be tuples of objects from more fundamental object sets, i.e. $O_i = (o_{i1}, ..., o_{ir}), o_{ij} \in \mathcal{O}_j$. However, the objects O_i are at the level of the mapping c.

iii.) The hypothesis map $\mathcal{C}(O^{(n)})$ is actually a composition of the maps $X : \mathcal{O}^{(n)} \to \mathcal{X}^n$ and $t : \mathcal{X}^n \to \mathcal{P}$ where $\mathcal{X}^n$ is a measurement space.

iv.) The pattern space $\mathcal{P}$ may be related to the data generating process or not. It is an *interpretation* space: a set of abstract, mutually exclusive *properties* which we wish to assign to object configurations.

We present some examples to clarify the pattern analysis framework.

Example 1.3.0.1 (Mean estimation)**.** We want to estimate the population mean height of swiss women given a sample of 100. We make no assumptions regarding the data generating process.

- $\mathcal{O} = \{\text{swiss women}\}$
- $O^{(100)}$ is the set of sampled women.
- $X(O^{(100)})$ are the heights of the sampled women.
- $\mathcal{X}^n \subset (\mathbb{R}^+)^{100}$ is the set of possible heights for the 100 women.
- $\mathcal{P} \subset \mathbb{R}^+$ is the set of possible population mean heights.

- $\mathcal{C}(O^{(100)}) = \mathcal{O}^{(100)} \times \mathcal{P}$ is the hypothesis class which does not depend on the size $n = 100$ of the object set.

Example 1.3.0.2 (Clustering - Population)**.** We want to cluster 100 people into 4 groups according to height and weight. We assume the underlying data generating process is a gaussian mixture with parameters $\{(\mu_1, \Sigma_1), ..., (\mu_4, \Sigma_4)\}$ with $\mu_i \in \mathbb{R}^2$ and $\Sigma_i \in \mathbb{R}^{2\times 2}$.

- $\mathcal{O} = \{\text{people}\}$.
- $O^{(100)}$ is the set of sampled people.
- $X(O^{(100)}) = \{(h_1, w_1), ..., (h_{100}, w_{100})\}$ are the heights and weights of the sampled people.
- $\mathcal{X}^n \subset \mathbb{R}^{100\times 2}$ is the set of possible heights and weights for the 100 people.
- $\mathcal{P} \subset \mathbb{R}^{2\times 4} \times \mathbb{R}^{2\times 2\times 4}$ is the set of possible population mean and covariances.
- $\mathcal{C}(O^{(100)}) = \mathcal{O}^{(100)} \times \mathcal{P}$ is the hypothesis class which does not depend on the size n of the object set.

Remark: Notice how our assumptions about the data generating process inform our choice of pattern space $\mathcal{P}$.

Example 1.3.0.3 (Clustering - Sample)**.** We want to cluster 100 people into 4 groups according to height and weight. We do not assume anything about the underlying data generating process and are just interested in finding a clustering that defines homogenous groups for *this* sample and not the entire population.

- $\mathcal{O} = \{\text{people}\}$.
- $O^{(100)}$ is the set of sampled people.
- $X(O^{(100)}) = \{(h_1, w_1), ..., (h_{100}, w_{100})\}$ are the heights and weights of the sampled people.
- $\mathcal{X}^n \subset \mathbb{R}^{100\times 2}$ is the set of possible heights and weights for the 100 people.
- $\mathcal{P} = \{1, 2, 3, 4\}^{100}$ are all the possible ways we can group 100 people into 4 groups.
- $\mathcal{C}(O^{(100)}) = \mathcal{O}^{(100)} \times \mathcal{P}$ is the hypothesis class which in this case **does depend** on the size n of the object set.

Example 1.3.0.4 (Dyadic data)**.** We are interested in predicting if a user will make a purchase at a given website, based on the age and gender of the person and on the type of website (there are m types). We don't assume anything about the data generating process but have already decided to model the probability of purchase using a logistic regression model.

- $\mathcal{O} = \mathcal{O}_p \times \mathcal{O}_w = \{\text{people}\} \times \{\text{websites}\}$.
- $O^{(n)} = \{(o_{1p}, o_{1w}), ..., (o_{np}, o_{nw})\}$ is the set of sampled person-website pairs.
- $X(O^{(n)}) = \{(a_1, g_1, t_1, p_1), ..., (a_n, g_n, t_n, p_n)\}$ are the age, geneder, website-type and purchase outcome of the sampled person-website pairs.

- $\mathcal{X}^n \subset \mathbb{R}^+ \times \{male, female\} \times \{1, ..., m\} \times \{0, 1\}$ is the sample space.
- $\mathcal{P} = \{(\beta_0, \beta_1, \beta_2) \in \mathbb{R}^3\}$ is the parameter space for the logistic model.
- $\mathcal{C}(O^{(n)}) = \mathcal{O}^{(n)} \times \mathcal{P}$ is the hypothesis class which does not depend on the size n of the object set.

As we can see the pattern analysis framework fits a wide range of problems. Although problems such as mean estimation and regression, in which hypothesis classes with infinite cardinality are involved, can be tackled using the pattern analysis framework, in the rest of this introductory chapter we focus on classes with a finite number of hypotheses. In other words we assume:

$$|\mathcal{C}(O^{(n)})| < \infty \tag{1.3.0.1}$$

1.4 Approximation sets

In classical statistical learning theory, in order to solve an inference decision problem, we choose a loss function $\rho(c, x)$ with which we construct the risk function $R(c) = \mathbb{E}_X[\rho(c, X)]$. We then choose a single hypothesis c^* that minimizes the empirical risk $\hat{R}(c, X^{(n)})$ for a given data set $X^{(n)}$:

$$c^*(X^{(n)}) \in \underset{c \in \mathcal{C}(O^{(n)})}{\arg\min}\ \hat{R}(c, X^{(n)}) \tag{1.4.0.1}$$

$$\hat{R}(c, X^{(n)}) = \frac{1}{n} \sum_{i=1}^{n} \rho(c, X_i) \tag{1.4.0.2}$$

If n is large then the *Emprirical Risk Minimizer* (ERM) will be close to the minimizer of the risk function $R(c, X^{(n)})$. However, in general we know that when n is not large then the ERM will tend to overfit the data. Instead of choosing a single hypothesis we can choose a subset of the hypothesis class which includes *good* hypotheses: hypothesis with low costs. Qualitatively, we would like this set to be composed of low cost hypotheses which we cannnot (partially) order further because their costs are statistically indistinguishable. The goal is to choose a subset of hypotheses that are stable with respect to fluctuations in the cost measurements. We may code this selection with a weight function $w_\beta(c, X^{(n)})$ over the hypothesis class where:

$$w : \mathcal{C} \times \mathcal{X}^n \times \mathbb{R}^+ \to [0, 1] \tag{1.4.0.3}$$

$$(c, X^{(n)}, \beta) \mapsto w_\beta(c, X^{(n)}) \tag{1.4.0.4}$$

$$w_\beta(c, X^{(n)}) = \begin{cases} 1, & \hat{R}(c, X^{(n)}) \leq \hat{R}(c^*, X^{(n)}) + 1/\beta \\ 0, & \text{otherwise} \end{cases} \tag{1.4.0.5}$$

Where β can be interpreted as the degree of certainty we have that c^* is the best solution. We may generalize the concept of approximation sets by allowing *fuzzy*, non-binary selection where hypotheses belong to the solution set to varying degrees, i.e. $w_\beta(c, X^{(n)}) \in [0, 1]$. In this case valid weight vectors satisfy:

$$w_\beta(c, X^{(n)}) \geq w_\beta(c', X^{(n)}) \Leftrightarrow \hat{R}(c, X^{(n)}) \leq \hat{R}(c', X^{(n)}) \tag{1.4.0.6}$$

The sum of the weights over the hypothesis class indicates the *equivalent* number of hypotheses selected. The bigger this sum the more unsure we are about c^*. We will sometimes say that $w_\beta(c, X^{(n)})$ is the *approximation set of hypotheses*, meaning that it encodes the (fuzzy) membership of the hypotheses in the set. A parametric family of weights which satisfies condition 1.4.0.6 is:

$$w_\beta^f(c, X^{(n)}) = \{w_\beta(c, X^{(n)}) = e^{-\beta f(\hat{R}(c, X^{(n)}))} : \beta \in \mathbb{R}^+, f \text{ increasing}\} \tag{1.4.0.7}$$

Notice that if we normalize the weights such that $\sum_{c \in \mathcal{C}(O^{(n)})} w_\beta(c, X^{(n)}) = 1$ we can interpret the weights as a posterior probability distribution over the hypothesis class.

In the classical statistical setting, when we have limited data, obtaining unbiased estimators often means these estimators have high variance: estimations change dramatically from one data set to the next. Lowering the variance can sometimes be achieved by introducing bias into our estimator. This is the bias-variance trade-off that, when there is limited data, is usually resolved through some sort of regularization. As we shall see in Section 1.6 the ASC approach leads to a resolution-stability trade-off which replaces the bias-variance trade-off. *Resolution* refers to the equivalent number of hypotheses selected while *stabiity* refers to obtaining similar approximation sets for different $X^{(n)} \in \mathcal{X}^n$. Adopting the ASC approach the trade-off becomes, do we obtain a very stable set of *good* hypotheses that don't change a lot depending on the data set but that is quite large (low resolution) or do we focus in on a small number of *very good* hypotheses but such that they will change from one data set to the next (unstable).

Notice that the β parameter in our weight function w_β is the resolution parameter that determines how this trade-off is resolved. Adopting the view of our normalized weight vector as a posterior over the hypothesis class, the higher β is the more probability is *spread* or *smoothed* among all the hypotheses. In limited data settings, choosing the parameter β corresponds to regularizing our empirical risk function.

In general, the justification for using the ASC approach is:

I **Inference**. It allows us to identify hypotheses which are similar in cost but which might be distinct according to other criteria not included in the cost function. This benefit is also common to bayesian inference.

II **Learnability**. For $\mathcal{C}(O^{(n)})$ to be *learnable*, ERM theory requires that it should not be too complex. In other words, $\mathcal{C}(O^{(n)})$ should have a finite *VC-dimension*. For certain problems, such as 1.3.0.3, the size of $\mathcal{C}(O^{(n)})$ increases too quickly in n, meaning that as $n \to \infty$ the empirical risk minimizer does not converge to the true risk minimizer. For this type of problem it is not even theoretically possible to converge to the true $c \in \mathcal{C}(O^{(n)})$ as $n \to \infty$ so an approximation set solution seems more reasonable.

1.5 Generalization capacity

Definition 1.5.0.1 (Boltzmann weights)**.** *If, from the parametric family 1.4.0.7, we choose $f(x) = x$ to construct our weight vector we obtain the so called Boltzmann weights:*

$$w_\beta(c, X^{(n)}) := e^{-\beta \hat{R}(c, X^{(n)})} \tag{1.5.0.1}$$

Definition 1.5.0.2 (Partition function)**.** *The sum over $\mathcal{C}(O^{(n)})$ of the Boltzmann weights is a function of the data $X^{(n)}$. We call it the partition function with respect to w_β and define it as:*

$$Z_\beta(X^{(n)}) := \sum_{c \in \mathcal{C}(O^{(n)})} w_\beta(c, X^{(n)}) \tag{1.5.0.2}$$

Definition 1.5.0.3 (Gibbs Distribution)**.** *The normalized Boltzmann weights define a Gibbs distribution, $P_G(c; \beta, X^{(n)})$ over $\mathcal{C}(O^{(n)})$ with respect to the cost function $\hat{R}(c, X^{(n)})$:*

$$P_G(c; \beta, X^{(n)}) := \frac{w_\beta(c, X^{(n)})}{Z_\beta(X^{(n)})} \tag{1.5.0.3}$$

This choice of weight vector can be justified from an information theoretic perspective. The Gibbs distribution is the maximum *entropy* distribution among all distributions $p(c)$ over $\mathcal{C}(O^{(n)})$ such that:

$$\mathbb{E}_{p(c)}[\hat{R}(c, X^{(n)})] = \mu_\beta \tag{1.5.0.4}$$

where μ_β is a non-increasing function of β. As we increase β, the resolution parameter, the expected cost with respect to the Gibbs distribution, decreases. In the limit, as $\beta \to \infty$, the Gibbs distribution becomes a single point mass distribution over c^* and $\mu_\beta \to \hat{R}(c^*, X^{(n)})$. The Gibbs distribution $P_G(c; \beta, X^{(n)})$ preserves the same (partial) ordering of $\mathcal{C}(O^{(n)})$ as $-\hat{R}(c, X^{(n)})$, but rescales so that differences in cost on the low end of the cost spectrum are *exaggerated* and differences in cost on the high end of the cost spectrum are *smoothed* out.

Figure 1.1 illustrates the mechanics of the smoothing of costs with the Boltzmann weight function. Costs scaled by the β parameter are plotted on the x-axis ($x = \beta \hat{R}(c, X^{(n)})$) and the corresponding Boltzmann weights on the y-axis ($y = w_\beta(c, X^{(n)}) = e^{-x}$). The black lines show the Boltzmann weights for two different resolution values: $\beta = 20, 1.5$. The green line represents the points $({}^{\log(\beta)}/_{\beta}, {}^{1}/_{\beta})$ which are the points that satisfy $\frac{dy}{dx} = -1$. Let the point of intersection between a given weight function $w_\beta(c, X^{(n)})$ indexed by β and the green line be called the critical point (x^c_β, y^c_β) for that β. For a given β, scaled costs to the right of the critical point x^c_β (high costs) are smoothed onto the interval $(0, y^c_\beta)$ while scaled costs to the left of x^c_β (low costs) are exaggerated onto the interval $(y^c_\beta, 1)$. This is how the Boltzmann weight function $w_\beta(c, X^{(n)})$ and the parameter β control the level of smoothing: for high resolution β only the lowest cost hypotheses remain relevant, while for low resolution levels most hypotheses retain some measure of relevance.

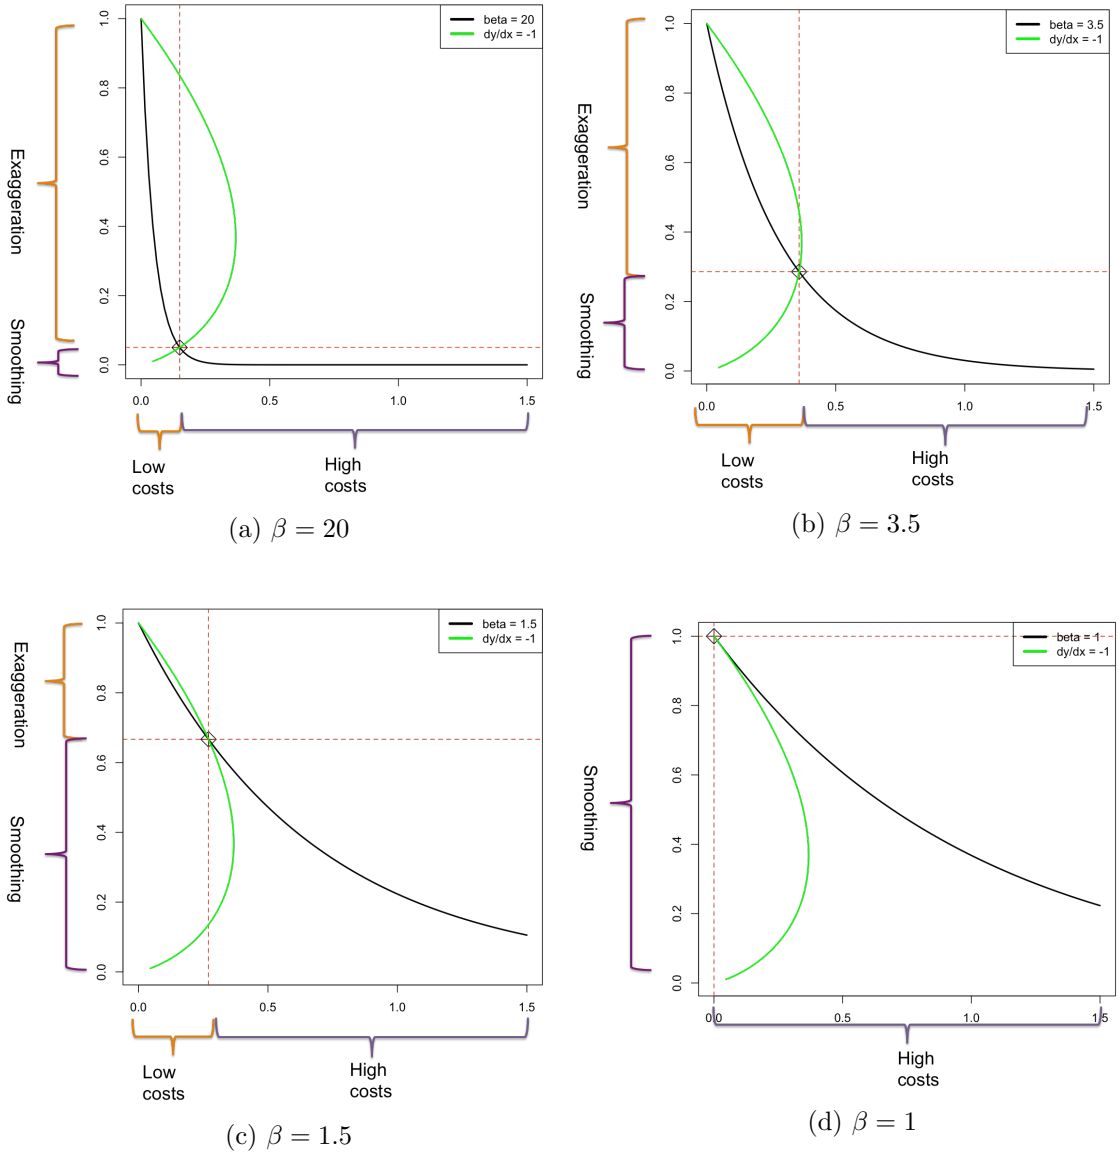

(a) $\beta = 20$ (b) $\beta = 3.5$

(c) $\beta = 1.5$ (d) $\beta = 1$

Figure 1.1: Boltzmann smoothing: smoothing of costs with Boltzmann weight function

The characteristics discussed above, are shared by all functions from the parametric family 1.4.0.7. These characteristics allow the Boltzmann weight function to be used in global optimization strategies such as *simulated* and *deterministic annealing* where the smoothing out of less important features in the cost surface in early iterations prevent the search algorithm from getting stuck in local minima. As is established in **?**, **?** and **?**, particular to the Gibbs distribution (for which $f(x) = x$), is the fact that for a given level of resolution, manifested as an expectation, $\mathbb{E}_{p(c)}[\hat{R}(c, X^{(n)})] = \mu_\beta$, that is a certain distance from $\hat{R}(c^*, X^{(n)})$, it has maximum entropy among distributions with this characteristic This means that if we use the Gibbs distribution to describe our uncertainty about the true hypothesis, the only information extracted from $X^{(n)}$ is that obtained using $\hat{R}(c, X^{(n)})$. Interpreting the entropy of a distribution as a measure of its uncertainty and supposing we know that $\mathbb{E}_{p(c)}[\hat{R}(c, X^{(n)})] = \mu_\beta$, then $P_G(c; \beta, X^{(n)})$ is the maximally non-comittal

distribution with respect to information different to that contained in this restriction. **?** established another characteristic that makes the Boltzmann weights and Gibbs distribution an appealing choice as the ASC weighting function: it is maximally stable. If we change our desired resolution level from β_1 to β_2, the change in the induced Gibbs distributions is minimal, in the L_2 norm sense, among any two distributions p_1 and p_2 that satisfy $\mathbb{E}_{p_1(c)}[\hat{R}(c, X^{(n)})] = \mu_{\beta_1}$ and $\mathbb{E}_{p_2(c)}[\hat{R}(c, X^{(n)})] = \mu_{\beta_2}$.

We have discussed the role of the resolution parameter β in the context of the resolution-stability trade-off, so how can we determine the best value of β? For this purpose **?** developed the concept of *Generalization Capacity* which we will first define and then describe in analogy to the *Channel Capacity* concept of information theory.

Definition 1.5.0.4 (Joint partition function)**.** *Before we define Generalization Capacity we define the joint partition function between two data sets which measures the equivalent number of hypotheses selected by a weighting function $w_\beta(c, X)$ for two different data sets X' and X'':*

$$\Delta Z_\beta(X', X'') = \sum_{c \in \mathcal{C}} w_\beta(c, X') w_\beta(c, X'') \tag{1.5.0.5}$$

Remarks:

i.) We have dropped the superindex (n) for better readability: $X' := X'^{(n)}$, $X'' := X''^{(n)}$, $O' := O'^{(n)}$ and $O'' := O''^{(n)}$.

ii.) As this definition already suggests GC will involve comparing the approximation sets obtained with different data sets of the same size.

iii.) For some pattern analysis problems such as 1.3.0.3 the hypothesis class depends on the object set $O^{(n)}$ so that $\mathcal{C}(O') \neq \mathcal{C}(O'')$. In this case we need a mapping $\psi : \mathcal{O}' \to \mathcal{O}''$ so that $\Delta Z_\beta(X', X'')$ can be properly defined as $\Delta Z_\beta(X', X'') = \sum_{c \in \mathcal{C}(O'')} w_\beta(c, \psi(X')) w_\beta(c, X'')$. Although these types of problems are very important in the context of ASC given the infinite VC dimension of the hypothesis class, for the sparse mean estimation problem described in Section 1.2 this is not the case, so we will simply assume, from now on, that $\mathcal{C} := \mathcal{C}(O') = \mathcal{C}(O'')$. This also means we can dispense with the mappings ψ in this exposition.

Definition 1.5.0.5 (Information Content)**.** *The information content retrievable from data (X', X'') by a cost function $\hat{R}(c, X)$ with resolution β is:*

$$I_\beta := \log \frac{|\mathcal{C}| \Delta Z_\beta(X', X'')}{Z_\beta(X') Z_\beta(X'')} \tag{1.5.0.6}$$

Remarks:

i. I_β is a normalized and rescaled version of $\Delta Z_\beta(X', X'')$ which measures the equivalent number of selected hypotheses with cost function $\hat{R}(c, X)$ for both data sets X' and X''.

ii. Since $\lim_{\beta \to 0} Z_\beta(X') = \lim_{\beta \to 0} Z_\beta(X'') = \lim_{\beta \to 0} \Delta Z_\beta(X', X'') = |\mathcal{C}|$ it holds that $\lim_{\beta \to 0} I_\beta = 0$, which means that for resolution $\beta = 0$, where all hypothesis are given a weight of 1, the information content is zero.

iii. Let $\mathcal{C} = \{c_1, ..., c_{|\mathcal{C}|}\}$, $a_i = w_\beta(c_i, X') \geq 0$ *and* $b_i = w_\beta(c_i, X'') \geq 0$, *then:*

$$0 \leq \frac{\Delta Z_\beta(X', X'')}{Z_\beta(X')Z_\beta(X'')} = \frac{\sum_{i=1}^{|\mathcal{C}|} a_i b_i}{(\sum_{i=1}^{|\mathcal{C}|} a_i)(\sum_{i=1}^{|\mathcal{C}|} b_i)} \tag{1.5.0.7}$$

$$= \frac{\sum_{i=1}^{|\mathcal{C}|} a_i b_i}{\sum_{i=1}^{|\mathcal{C}|}\sum_{j=1}^{|\mathcal{C}|} a_i b_j} = \frac{\sum_{i=1}^{|\mathcal{C}|} a_i b_i}{\sum_{i=1}^{|\mathcal{C}|} a_i b_i + \sum_{i\neq j} a_i b_j} \leq 1 \tag{1.5.0.8}$$

and we can see that $\lim_{\beta\to\infty} I_\beta \leq \log|\mathcal{C}|$. *This means that for maximum resolution the information content can reach up to the log-size of the hypothesis class.*

Definition 1.5.0.6 (Generalization Capacity)**.** *The generalization capacity of a cost function* $\hat{R}(c, X)$ *defined over a hypothesis class* $\mathcal{C}$ *and data space* $\mathcal{X}^n$ *is:*

$$I := \max_{\beta\in\mathbb{R}^+} \mathbb{E}_{(X',X'')} I_\beta \tag{1.5.0.9}$$

Remarks:

i. Since $\lim_{\beta\to 0} I_\beta = 0$ *we have that* $\lim_{\beta\to 0} I = 0$ *and,s*

ii. since $\lim_{\beta\to\infty} I_\beta \leq \log|\mathcal{C}|$ *we have that* $\lim_{\beta\to\infty} I \leq \log|\mathcal{C}|$

1.6 GC and Shannon's noisy-channel coding theorem

To motivate the relevance of the Generalization Capacity as an important quantity in itself aswell as a criterion for deciding between cost functions in a pattern analysis problem, we briefly study Shannon's Noisy-Channel Coding Theorem, the communication protocol suggested therein and the role of *Channel Capacity*. We then move from the communication context to the pattern learning context and study an analagous *learning protocol* suggested by **?** where the generalization capacity emerges as a natural counterpart to channel capacity. The exposition of Shannon's Noisy-Channel Coding Theorem is based on **?** and **?**.

Shannon's Noisy-Channel Coding Theorem deals with the rate at which information can be passed through a channel so we first define what information and channels are.

Definition 1.6.0.1 (Shannon Information)**.** *The Shannon information of an outcome* x *of a random variable* $X \in \mathcal{X}$, *where* $\mathcal{X}$ *is a finite set and* $p(x)$ *is the probability distribution of* X *is:*

$$I(x) = -\log p(x)$$

If we interpret the informativeness of an outcome in terms of the worth of knowing its value the following properties make this a useful measure of information:

i.) $I(x) \geq 0 \ \forall x$

ii.) *Assigns 0 to a certain outcome*

iii.) *The rarer an outcome the more informative:* $p(x) = p, p(y) = q, p < q \Rightarrow I(x) > I(y)$

iv.) *It is continuous in* $p(x)$*:* $\forall \epsilon > 0 \ \exists \ \delta > 0 : \ $ *if* $|p - q| < \delta \Rightarrow |I(x) - I(y)| < \epsilon$

v.) *Additivity:* $p_{XY}(x, y) = p_X(x)p_Y(y) \Rightarrow I(x, y) = I(x) + I(y)$

Definition 1.6.0.2 (Entropy)**.** *The entropy of a random variable* $X \in \mathcal{X}$*, where* $\mathcal{X}$ *is a finite set, is its expected Shannon Information:*

$$H(X) := -\mathbb{E}_X[\log p(X)] = \mathbb{E}_X[I(X)] = \sum_{x \in \mathcal{X}} p(x) \log p(x) \tag{1.6.0.1}$$

We can interpret entropy as the average information rate of a random variable. If we want to send *messages* from a finite message set $\mathcal{W} = \{1, ..., M\}$, we may define the information rate of the message set by assuming messages will be sent according to the uniform distribution. In this case:

$$H(\mathcal{W}) := -\sum_{i=1}^{M} \frac{1}{M} \log \frac{1}{M} = \log M \tag{1.6.0.2}$$

We now define joint and conditional entropy.

Definition 1.6.0.3 (Joint Entropy)**.** *For random variables* $X \in \mathcal{X}$ *and* $Y \in \mathcal{Y}$*, with* $\mathcal{X}$ *and* $\mathcal{Y}$ *finite sets, the joint entropy of* X *and* Y *is defined as:*

$$H(X, Y) := -\sum_{(x,y) \in \mathcal{X} \times \mathcal{Y}} p(x, y) \log p(x, y) = -\mathbb{E}_{(X,Y)} \log p(X, Y) \tag{1.6.0.3}$$

Definition 1.6.0.4 (Conditional Entropy)**.** *For random variables* $X \in \mathcal{X}$ *and* $Y \in \mathcal{Y}$*, with* $\mathcal{X}$ *and* $\mathcal{Y}$ *finite sets, the joint entropy of* Y *given* X *is defined as:*

$$H(Y|X) := -\sum_{(x,y) \in \mathcal{X} \times \mathcal{Y}} p(x, y) \log p(y|x) = -\mathbb{E}_{(X,Y)} \log p(Y|X) \tag{1.6.0.4}$$

Conditional entropy is a measure of the mean information left in Y once we know the outcome of X. It turns out that $H(X, Y) = H(X) + H(Y|X)$ so the joint entropy can be interpreted as the mean amount of information in X plus the mean amount of information left in Y once the outcome of X is known.

Definition 1.6.0.5 (Mutual Information)**.** *For random variables* $X \in \mathcal{X}$ *and* $Y \in \mathcal{Y}$*, with* $\mathcal{X}$ *and* $\mathcal{Y}$ *finite sets, the mutual information between* X *and* Y *is defined as:*

$$I(X; Y) := \sum_{(x,y) \in \mathcal{X} \times \mathcal{Y}} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} = \mathbb{E}_{(X,Y)} \log \frac{p(X, Y)}{p(X)p(Y)} \tag{1.6.0.5}$$

Using that $I(X;Y) = H(X) - H(X|Y)$ we can interpret the mutual information as the reduction in information left in X once Y is known (or vice versa). Alternatively, we can interpret $I(X;Y)$ as the information that is common to X and Y. If X and Y are independent then they have no information in common and if X depends deterministically on Y then they have the same information.

Definition 1.6.0.6 (Discrete Channel)**.** *Let $\mathcal{X}$ and $\mathcal{Y}$ be discrete sets and $p(y|x)$ be a transition matrix from $\mathcal{X}$ to $\mathcal{Y}$ that is a valid distribution for all x. Then the tuple $(\mathcal{X}, p(y|x), \mathcal{Y})$ is a discrete channel where $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ are the input and output respectively.* ***Remark****: We sometimes refer to the discrete channel simply as $p(y|x)$.*

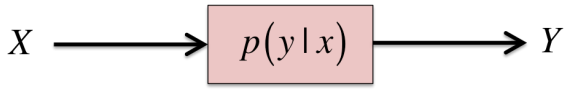

Figure 1.2: Discrete channel: schematic of input and output of a discrete channel

If Input $X = x$ is *sent* through the channel then the output Y is distributed according to $p(y|x)$.

Definition 1.6.0.7 (Discrete Memoryless Channel (DMC))**.** *A discrete memoryless channel is a discrete channel $(\mathcal{X}, p(y|x), \mathcal{Y})$ such that if a sequence of inputs $X_1, X_2, ...$ are sent through the channel then:*

$$Y_t \perp\!\!\!\perp \{X_1, Y_1, X_2, Y_2, ..., X_{t-1}, Y_{t-1}\}|X_t \tag{1.6.0.6}$$

Intuitively, the channel *forgets* all previous communication such that the ouptut of Y_t only depends on the input X_t and on the distribution $p(y|x)$.

Definition 1.6.0.8 (Capacity of a DMC)**.** *The capacity of a DMC $(\mathcal{X}, p(y|x), \mathcal{Y})$ is defined as:*

$$C := \max_{p(x)} I(X;Y) \tag{1.6.0.7}$$

Where X and Y are the input and output of the channel.

The capacity of a DMC is a measure of the amount of common information between the input and output in the most optimistic scenario. As we will see later in this section, Shannon's Noisy-Channel Coding Theorem shows why the capacity of a DMC is an important quantity.

Since a given channel $(\mathcal{X}, p(y|x), \mathcal{Y})$ only takes as input $X \in \mathcal{X}$ we need an encoder function to transform our message $w \in \mathcal{W}$ into an acceptable input. If $|\mathcal{X}| = M$ then we may simply assign each message an element of the input set $\mathcal{X}$, however this doesn't help us avoid errors in communication. If the channel transforms the message such that the output is not the same as the input then an error will occur.

If $|\mathcal{X}| > M$ we have some *slack* in our input set $\mathcal{X}$ which may help us to avoid errors. Suppose that $\mathcal{X} = \{x_1, x_2, x_3\}$ and $\mathcal{W} = \{0, 1\}$ then we can assign 0 to x_1 and 1 to x_2 and x_3. If an x_2 is sent through and the channel distorts it into an x_3 we still avoid error.

If $|\mathcal{X}| \leq M$ we may add *slack* to our coding scheme by encoding each message $w \in \mathcal{W}$ with a sequence of n symbols $x_i \in \mathcal{X}$. In this case we have $|\mathcal{X}|^n$ sequences to encode M messages. If we let n grow then we increase the slack in our code and so reduce the probability of error, especially if we assign sets of sequences to each message in a *smart* way. To prove the Noisy-Channel Coding theorem Shannon constructed such a smart assignment procedure using ideas of *typicality* which we explore somewhat further on.

We can already touch on how the pattern analysis problem bares some resemblance to the problem of sending a message through a noisy channel: in the former there is some *truth* or property in nature which is a hypothesis $c \in \mathcal{C}$ and it is encoded in a slack way by a data generating mechanism such that for each hypothesis there correspond many possible data sets $X^{(n)} \in \mathcal{X}^n$.

Having broached the idea of slack codes we now define an (n, M) code and give a schematic description of Shannon's communication scenario.

Definition 1.6.0.9 ((n, M) code)**.** *An (n, M) code for a DMC $(\mathcal{X}, p(y|x), \mathcal{Y})$ is defined by an encoding function f and a decoding function g:*

$$f : \mathcal{W} \to \mathcal{X}^n \qquad\qquad g : \mathcal{Y}^n \to \mathcal{W}$$

Where:

- $\mathcal{W} = \{1, ..., M\}$ *is called the message set,*
- $f(1), ..., f(M) \in \mathcal{X}^n$ *are the codewords and*
- $\{f(1), ..., f(M)\}$ *is the codebook.*

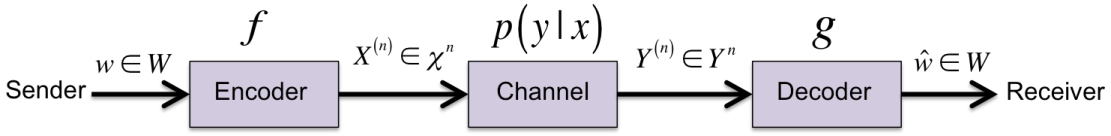

Figure 1.3: Communication channel: schematic of communication channel

The rate of an (n, M) code is defined as $\frac{1}{n} \log M$. This corresponds to the rate of a uniform random variable over the message set $\mathcal{W}$ divided by n so that it is in the units of bits (or nats depending on the base of the logarithm) per symbol $X \in \mathcal{X}$ and not bits per sequence $X^{(n)} \in \mathcal{X}^n$.

A rate R for a DMC is said to be *asymptotically achievable* if there exists an (n, M) code that for a sufficiently large n can transmit at a rate arbitrarily close to R with arbitrary precision.

Theorem 1.6.0.10 (Shannon's Noisy-Channel Coding Theorem)**.** *A rate R is asymptotically achievable for a DMC $\iff$ $R \leq C$*

This theorem justifies the *Capacity* of a channel as an interesting quantity: it implies that we can achieve error-free communication at a rate equal to the capacity of the DMC. The proof of this part of the theorem involves proposing the (n, M) code shown below and then proving that for an n such that the rate $R = \frac{1}{n} log M$ of the code is close to the capacity of the DMC, the probability of error $P(w \neq \hat{w})$ is small. In this work we are especially interested in the (n, M) code proposed in Shannon's proof since it forms the basis of a

similar coding scheme and communication protocol in which generalization capacity plays an analogous role to that of channel capacity.

Definition 1.6.0.11 (Shannon's (n,M) code)**.** *The (n, M) code proposed is the following:*

1 Sample M sequences $X^{(n)}$ uniformly at random from $\mathcal{X}^n$ and randomly assign each sequence sampled to one of the messages. This establishes the encoding function f. Both sender and receiver have the codebook $\{f(1), ..., f(M)\}$.

2 Compare the joint entropy of $H(X, Y)$, to the empirical entropy of the the pairs of sequences $(f(1), Y^{(n)}), ..., (f(M), Y^{(n)})$. Choose message i such that the empirical entropy of $(f(i), Y^{(n)})$ is close to the entropy $H(X, Y)$. If there is more than one pair of sequences that satisfies this condition decode to $\hat{w} = 1$. This establishes the decoding function g.

The proof that this (n, M) code can asymptotically achieve a rate $R = C$ involves the concept of *typicality* which is an application of the *Law of Large Numbers*. Although we do not give the formal proof we give a sequential illustration of the ideas.

i.) **Start with a large n for a slack code.** First we choose n large so that we have a lot of slackness in our code, more than we will need, and uniformly at random choose M sequences $X^{(n)} \in \mathcal{X}^n$ as our codewords.

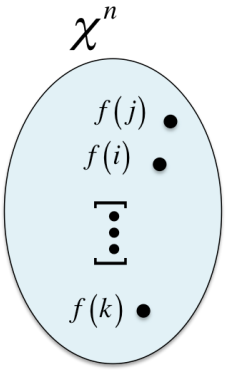

Figure 1.4: Shannon code: creation of codebook

ii.) **Channel sends messages to non-overlapping regions**. We have chosen n so large that even with a lot of noise, when $f(1), ..., f(M) \in \mathcal{X}^n$ are transformed into $Y_1^{(n)}, ..., Y_M^{(n)} \in \mathcal{Y}^n$ by the channel $(\mathcal{X}, p(y|x), \mathcal{Y})$ the probability that Y_i and Y_j are *close* for any $i \neq j$ is essentially zero. The i-th region $\mathcal{R}_i$ represents the sequences $Y^{(n)} \in \mathcal{Y}^n$ which are jointly typical with $f(i)$. Given that the sequence passed over the channel is $f(i)$, the probability that the sequence received by the decoder is outside this region is essentially zero. If $Y_j^{(n)} \in \mathcal{R}_i$ and $\nexists k \neq i : Y_j^{(n)} \in \mathcal{R}_k$ then $g(Y_j^{(n)}) = i$.

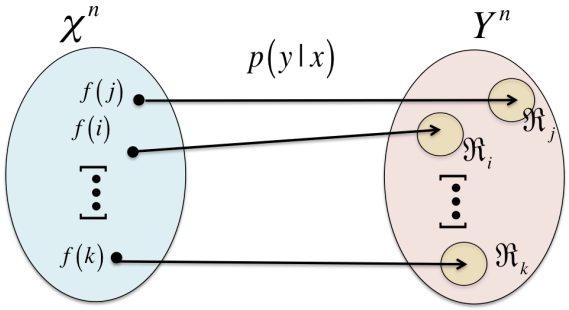

Figure 1.5: Shannon code: decoding

iii.) **Decrease n until $\mathcal{Y}^n$ is tight around regions**. Since n is large the rate $R = \frac{1}{n}logM$ of the code is low. Shannon's Noisy-Channel Coding Theorem says that we can decrease n so that the rate increases to close to C and the error stays very small. By decreasing n we decrease the size of $\mathcal{Y}^n$ so that all the regions $\mathcal{R}_i$ are tightly crowded within. If we are at capacity, the overlap between the regions is still essentially zero, but if we make $\mathcal{Y}^n$ any smaller by decreasing n further, the overlap will start to grow, meaning the probability of error grows.

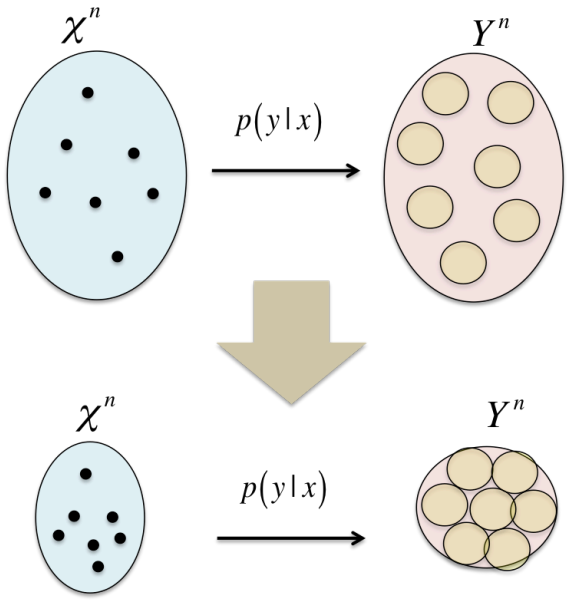

Figure 1.6: Shannon code: towards capacity

As we have already hinted at, the pattern analysis problem can be seen as a special case of the communication problem. We explore this further by proposing the following *Idealized Learning Protocol* :

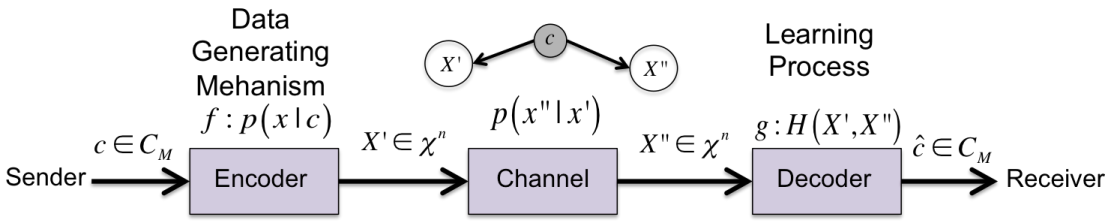

Figure 1.7: Idealized learning protocol

The communication protocol has the following characteristics:

1 The sender picks M uniformly at random to construct $\mathcal{C}_M = \{c_1, ..., c_M\} \subseteq \mathcal{C}$: $c_i \sim p(c) = \frac{1}{|\mathcal{C}|}$.

2 The sender selects a message $\widetilde{c}_s \in \mathcal{C}_M$ uniformly at random: $\widetilde{c}_s \sim p_s(c) = \frac{1}{M}$. **Remark:** we use the tilde to separate $c_i \sim p(c) = \frac{1}{|\mathcal{C}|}$, $i \in \{1, ..., M\}$ from $\widetilde{c}_s \sim p_s(c) = \frac{1}{M}$.

3 The sender and receiver have access to the data generating mechanism $p(x|c)$ which they use to set up the following (n, M) code:

   a The encoder function f is constructed by randomly sampling M times a data set of size n from the data generating mechanism $p(x|c)$ to obtain the codebook $\{f(c_1), ..., f(c_M)\}$ where $f(c_i) \in \mathcal{X}^n$. Notice that in many problems such as clustering and regression $\mathcal{X} = \mathbb{R}^d$ so that $\mathcal{X}^n = \mathbb{R}^{n \times d}$, i.e. we code each message using a data matrix.

   b Our decoding function g works as in Shannon's (n, M) code except that $\mathcal{X}$ is not necessarily a finite set so we might need to use joint *differential entropy* instead of joint entropy.

4 We set up a channel $(\mathcal{X}^n, p(x''|x'), \mathcal{X}^n)$ with our knowledge of the hypothesis and data generating mechanisms $p_s(c)$ and $p(x|c)$ respectively:

$$p(x''|x') \propto p(x', x'') = \sum_{c \in \mathcal{C}} p(c, x', x'') = \sum_{c \in \mathcal{C}} p_s(c) p(x', x''|c) \tag{1.6.0.8}$$

$$= \frac{1}{M} \sum_{c \in \mathcal{C}} p(x'|c) p(x''|c) \tag{1.6.0.9}$$

   Notice that we are assuming that successive sample sets of size n from the data generating mechanism are independent given c.

With the exception that we are using what may be a set $\mathcal{X}$ with infinite cardinality to code the message set $\mathcal{C}_M = \{c_1, ..., c_M\}$ the above Idealized Learning Protocol corresponds to the previous communication protocol. The channel $p(x''|x')$ characterizes the noisiness of the pattern analysis problem since in a noise-free scenario we would obtain the same data set for each realization of the data generating mechanism, i.e. $X' = X''$. Since messages are always selected uniformly at random according to $p_s(x') = p_s(c) = \frac{1}{M}$ we may consider the capacity of the channel to be $I(X'; X'')$. The capacity of the channel is a measure of the noisiness (the higher the capacity the less is the noise) and is an upper bound on the rate at which any learning algorithm can extract information from specific realizations $X \in \mathcal{X}^n$, that generalizes accross realizations, i.e. information about c and not about

the noise. In the above Idealized Learning Protocol we can achieve the capacity rate, as before, by choosing a suitable n. Since this scenario is highly idealized we make successive changes to it until we arrive at the more useful Learning Protocol proposed by **?** and from which generalized capacity is derived:

I **Change *expressiveness* of codebook instead of size of code sequences**. Suppose we can no longer change n, the size of our sequence $X \in \mathcal{X}^n$, i.e. it is fixed. Instead we are allowed to change M the number of selected to form $\mathcal{C}_M$. We can now achieve a rate close to capacity by increasing M instead of decreasing n. Our message set $\mathcal{C}_M$ and codebook $\{f(c_1), ..., f(c_M)\}$ become more *expressive* as we increase M.

II **Coding based on transformation set**. Additionally, suppose we can only use the data generating mechanism $p(x|c)$ and channel $p(x''|x')$ once. We still know the form of $p(x|c)$ and $p(x''|x')$ (and so can calculate entropies for decoding) but can only generate with it once. Furthermore, suppose that we don't know what hypothesis c_D is selected and passed to the data generating mechanism $p(x|c)$. Since we don't have any information about c_D other than the data X' and X'', we use a uniform prior $c_D \sim p(c) = \frac{1}{|\mathcal{C}|}$ to describe our uncertainty regarding the true c_D. All this means we can only generate two data sets $X', X'' \in \mathcal{X}^n$: we generate X' using $p(x|c)$ and then send it through the channel $p(x'|x'')$ to get X''. Without access to the data generating mechanism and the true c_D we need an alternative way to construct a codebook (i.e. another f for our (n, M) code). Consider the set of unique maps:

$$\mathcal{T}^h = \{t^h \neq u^h \in \mathcal{T} : \forall c \neq d \in \mathcal{C} \Rightarrow t^h(c) \neq t^h(d) \text{ and } t^h(c) \neq u^h(c)\} \tag{1.6.0.10}$$

Where $\mathcal{T} = \{t : \mathcal{C} \to \mathcal{C}\}$.

Notice the following properties about $\mathcal{T}^h$:

(a) $|\mathcal{T}^h| = |\mathcal{C}|$

(b) If you apply a fixed $t^h \in \mathcal{T}^h$ on all $c \in \mathcal{C}$ you get C again.

(c) If you apply all $t^h \in \mathcal{T}^h$ on a fixed $c \in \mathcal{C}$ you get C again.

Additionally, consider the set of maps:

$$\mathcal{T}^D = \{t^D : \mathcal{X}^n \to \mathcal{X}^n\} \tag{1.6.0.11}$$

We assume that we have a mapping $\phi : \mathcal{T}^h \to \mathcal{T}^D$ such that for a given $t^h \in \mathcal{T}^h$, $\phi(t^h) = t^D \in \mathcal{T}^D$ and:

$$p(c|x) = p(t^h(c)|t^D(x)) \tag{1.6.0.12}$$

Where the posteror $p(c|x)$ is obtained from the data generating mechanism $p(x|c)$ and the prior $p(c)$:

$$p(c|x) \propto p(x|c)p(c) \tag{1.6.0.13}$$

This means that if data set X is generated under hypothesis c with the data generating mechanism $p(x|c)$ then, for a given *move* t^h within the hypothesis space $\mathcal{C}$ we know how to make a corresponding *move* t^D in the coding/data space $\mathcal{X}^n$. The assumption that we can obtain a mapping ϕ is reasonable in some contexts such as in the sparse mean estimation problem that is the main topic of this work. In other pattern analysis problems such as in the mean estimation problem 1.3.0.1 where the hypothesis class $\mathcal{C}$ has infinite cardinality, the validity of this assumption is not clear.

With the above assumption we will be able to encode M $c_i \in \mathcal{C}_M$ into a codebook $\{f(c_1), ..., f(c_M)\} \subseteq \mathcal{X}^n$ However we can only pass one data set through the channel $p(x''|x')$ and arbitrarily choose to pass $f(c_D) = X' = x'_0$ which gives the random output $X'' = x''_0$. Observe that:

$$\begin{aligned}
p(t_i^D(x'), t_i^D(x'')) &= \sum_{c\in\mathcal{C}} p(c, t_i^D(x'), t_i^D(x'')) && (1.6.0.14)\\
&= \sum_{c\in\mathcal{C}} p(c)p(t_i^D(x'), t_i^D(x'')|c) && (1.6.0.15)\\
&= \sum_{c\in\mathcal{C}} p(c)p(t_i^D(x')|c)p(t_i^D(x'')|c) && (1.6.0.16)\\
&= \sum_{c\in\mathcal{C}} p(t_i^h(c))p(t_i^D(x')|t_i^h(c))p(t_i^D(x'')|t_i^h(c)) && (1.6.0.17)\\
&= \sum_{c\in\mathcal{C}} p(c)p(x'|c)p(x''|c) && (1.6.0.18)\\
&= \sum_{c\in\mathcal{C}} p(c)p(x', x''|c) && (1.6.0.19)\\
&= \sum_{c\in\mathcal{C}} p(c, x', x'') = p(x', x'') && (1.6.0.20)
\end{aligned}$$

Where we have used assumption 1.6.0.12 and the properties of 1.6.0.10. This implies that:

$$p(t_i^D(x'')|t_i^D(x')) = p(x''|x') \tag{1.6.0.21}$$

We wish to *mimic* the channel $p(x''|x')$ by mimicking the noise process that *contaminated* x'_0 to produce x''_0. We may think of the output of the channel $p(x''|x')$ as a function α of the input x' and a noise realization of some N. We then have that

$$\begin{aligned}
X'' = \alpha(x'_0, N) &\sim p(x''|x'_0) && (1.6.0.22)\\
t_i^D(X'') = \alpha(t_i^D(x'_0), N) &\sim p(x''|x'_0) && (1.6.0.23)
\end{aligned}$$

Since we have passed x'_0 through the channel and have observed $X'' = x''_0$ we implicitly have a noise observation $N = n_0$. Although, we can't pass $t_i^D(x'_0)$ through the channel

to get $t_i^D(X'')$ (recall we are only allowed to use the channel once, and we have already used it to pass x_0' through), we can mimic the output with $t_i^D(x_0'')$ which is the result of evaluating α on n_0 instead of on a new realization of N:

$$t_i^D(X'') = \alpha(t_i^D(x_0'), N) \tag{1.6.0.24}$$
$$t_i^D(x_0'') = \alpha(t_i^D(x_0'), n_0) \tag{1.6.0.25}$$

With the set $\mathcal{T}^c$ and the mapping ϕ we have the necessary elements to replace our encoding function f. Incorporating the changes to the encoding function and channel, the learning protocol, thus far, consists of:

1 The sender picks M transformations uniformly at random to construct $\mathcal{T}_M^h = \{t_1^h, ..., t_M^h\} \subseteq \mathcal{T}^h$: $t_i^h \sim p(t^h) = \frac{1}{|\mathcal{T}^h|} = \frac{1}{|\mathcal{C}|}$. Even though we don't know what c_D is we can set up our message set C_M using $\mathcal{T}_M^h$: $C_M = \{c_1, ..., c_M\} = \{t_1^h(c_D), ..., t_M^h(c_D)\}$. In fact, we may now say that $\mathcal{T}_M^h$ *is* the message set.

2 The sender selects a message $\widetilde{t}_s^h \in \mathcal{T}_M^h$ uniformly at random: $\widetilde{t}_s^h \sim p_s(t^h) = \frac{1}{M}$. **Remark:** we use the tilde to separate $t_i^h \sim p(t^h) = \frac{1}{|\mathcal{C}|}$, $i \in \{1, ..., M\}$ from $\widetilde{t}_s^h \sim p_s(t^h) = \frac{1}{M}$.

3 The sender and receiver have access to the transformation set $\mathcal{T}_M^h$ and the mapping ϕ which they use to set up the following (n, M) code:

   a The encoder function f is constructed by applying ϕ to each $t_i^h \in \mathcal{T}_M^h$ to construct $\mathcal{T}_M^D = \{t_1^D, ..., t_M^D\} = \{\phi(t_1^h), ..., \phi(t_M^h)\}$. Our codebook vector is then $\{f(c_1), ..., f(c_M)\} = \{f(t_1^h(c_D)), ..., f(t_M^h(c_D))\} = \{t_1^D(X'), ..., t_M^D(X')\}$.

   b Since we still know the distribution $p(x|c)$ we may use it to calculate $p(x', x'') = \frac{1}{M}\sum_{c\in\mathcal{C}} p(x''|c)p(x''|c)$. Notice that since we know how $\widetilde{c}_s = \widetilde{t}_s^h(c_D)$ is selected we use $p_s(c) = \frac{1}{M}$. With $p(x', x'')$ we can calculate the joint entropy $H(X', X'')$ and use it for decoding as before.

4 We can only send one data set $X' = x_0'$ through the channel $(\mathcal{X}^n, p(x''|x'), \mathcal{X}^n)$ and so can only observe one output $X'' = x_0''$. However, we may mimic the behavior of the channel for other input data sets $t_i^D(X')$ by using, as derived above, that:

$$p(t_i^D(x'')|t_i^D(x')) = p(x''|x') \tag{1.6.0.26}$$

This means that we may mimic the channel $p(x''|x')$ by outputing $t_i^D(x_0'')$ for a given input $t_i^D(x_0')$. This output corresponds to the output the actual channel $(\mathcal{X}^n, p(x''|x'), \mathcal{X}^n)$ would have given for an input $t_i^D(x_0')$, assuming the same realization of the noise process as ocurred when X' passed through the channel. Since the noise realization for one data set, is made up of n components, the hope is that the realization observed summarizes the noisiness of the channel. In other words we hope that applying this noise realization to any input data set $t_i^D(x_0')$, the output data set $t_i^D(x_0'')$ is similar to that we would get by passing $t_i^D(X')$ through the real channel (to which we no longer have access).

The assumption that we still know the distribution $p(x|c)$ is the last *idealized*, unrealistic element of our learning protocol. The last change to the protocol involves dispensing with this assumption. To do so we will use the Boltzmann approximation sets discussed in sections 1.4 and 1.5.

III **Decoding based on approximation sets**. Finally suppose we don't know the distribution of the data generating mechanism $p(x|c)$ or that of the channel $p(x''|x')$. This means we need a new decoding function g since we don't know $p(x', x'')$ and so cannot use the joint entropy $H(X', X'')$ for decoding. This is where our learning algorithm comes to the fore in the form of the cost function $\hat{R}(c, X)$ and the Gibbs distributions corresponding to both data sets: $P_G(c; \beta, X')$ and $P_G(c; \beta, X'')$. We first describe the new decoding function g and then discuss the ideas behind it and its relationship to genealization capacity.

Let

$$\Delta Z_\beta^j := \sum_{c \in \mathcal{C}} w_\beta(c, t_j^D(X')) w_\beta(c, t_s^D(X'')) \tag{1.6.0.27}$$

Then the decoding rule is

$$g_\beta(X'') \in \underset{j \in \{1,\ldots,M\}}{\arg\max}\ \Delta Z_\beta^j \tag{1.6.0.28}$$

Where ties are resolved by taking the minimum j. Before discussing how to choose the resolution parameter β we can show the final Learning Protocol schematic.

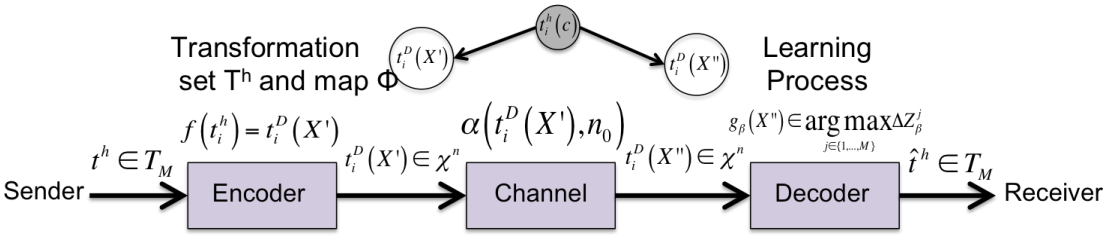

Figure 1.8: Learning protocol

What β should we use? In general we can use any β, however to find the channel capacity we must choose it so that for a given M (which determines the rate $R = \frac{1}{n} \log M$, given n is fixed) we can achieve error-free communication. We then increase M to M^* such that if we increase it any further there exists no β that allows error-free communication. This M^* determines the maximum achievable rate of our code. We illustrate the process of finding the maximum achievable rate of the above (n, M) code.

i **Start with low expressiveness M for slack code**. We uniformly sample M $c_i \in \mathcal{C}_M$ by sampling M transformations $t_i^h \in \mathcal{T}_M^h$ and use the corresponding set

$\mathcal{T}_M^D$ to build our codebook. We select a transformation $\widetilde{t}_s^h$ and pass the econcoded message $\widetilde{t}_s^D(X')$ through the channel to decoder that receives $\widetilde{t}_s^D(X'')$.

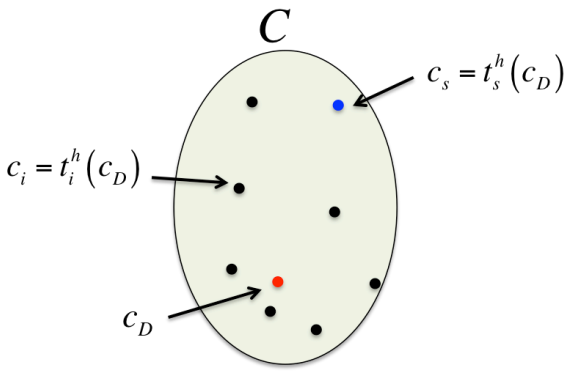

Figure 1.9: Learning protocol: creation of codebook

ii **Create high resolution β approximation sets**. Using high resolution β, we calculate M approximation sets $w_\beta(c, t_j^D(X'))$, one for each codeword in the codebook. Using the received data set, we calculate an additional approximation set $w_\beta(c, \widetilde{t}_s^D(X''))$, and apply decoding rule. In the case illustrated in figure 1.10 $\Delta Z_\beta^j = 0\ \forall j$ so by default we decode message to $t_1^h(c_D)$. The red circles represent the approximation sets $w_\beta(c, t_j^D(X'))$ while the blue circle represents the approximation set $w_\beta(c, \widetilde{t}_s^D(X''))$. Recall from Section 1.4 that although, strictly speaking, $w_\beta(c, X)$ is a weight vector over the entire hypothesis class $\mathcal{C}$, the circles represent the subset of $\mathcal{C}$ where the *majority* of the weight is supported.

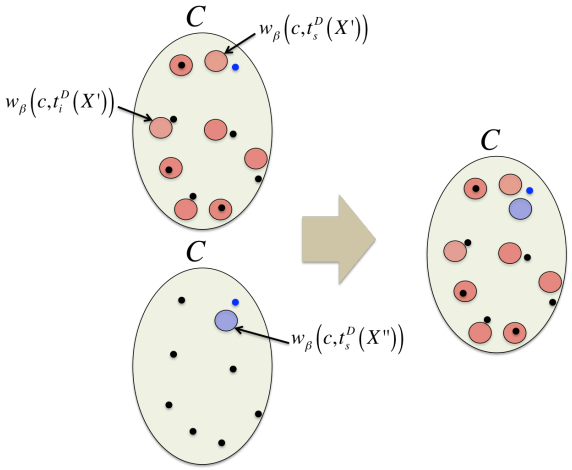

Figure 1.10: Learning protocol: decoding with high resolution

iii **Lower resolution β to increase intersection**. Since the current resolution level doesn't allow significant intersection between $w_\beta(c, \widetilde{t}_s^D(X''))$ and any of the $w_\beta(c, t_j^D(X'))$ approximation sets we lower the resolution level until there is some intersection meaning we can decode the message without error. Since the hypothesis class $\mathcal{C}$ is not *cluttered* with approximation sets, because we are not near capacity, the intersection need not represent a large percentage of the respective

approximation sets.

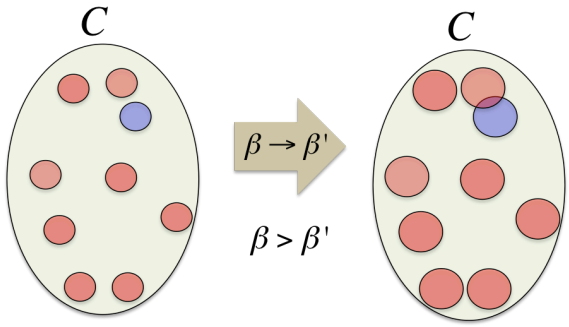

Figure 1.11: Learning protocol: decoding with low resolution

iv **Increase expressiveness M, while adjusting resolution β.** In finding the maximum rate of our code we increase the size of the message set which means more approximation sets over $\mathcal{C}$. However, as happens in part (b) of figure 1.12, for a given resolution β the approximation set $w_\beta(c, \widetilde{t}_s^D(X''))$ intersects with more than one approximation set $w_\beta(c, t_j^D(X'))$ meaning the probability of decoding the wrong message increases. We can, as is shown on part (c), fix this by decreasing the resolution but this increases the size of the approximation sets $w_\beta(c, t_j^D(X'))$ so that they intersect between themselves. This also increases the probability of error since the set $w_\beta(c, \widetilde{t}_s^D(X''))$ intersects with several $w_\beta(c, t_j^D(X'))$ sets simultaneously. Finally in part (d) we obtain obtain the maximum rate of our code: if we add any more approximation sets $w_\beta(c, t_j^D(X'))$, they will become cluttered between themselves (we go back to the situation in part c) and if we then fix this by increasing resolution, more than one of these sets will intersect significantly with $w_\beta(c, \widetilde{t}_s^D(X''))$ (we go back to the situation in part b).

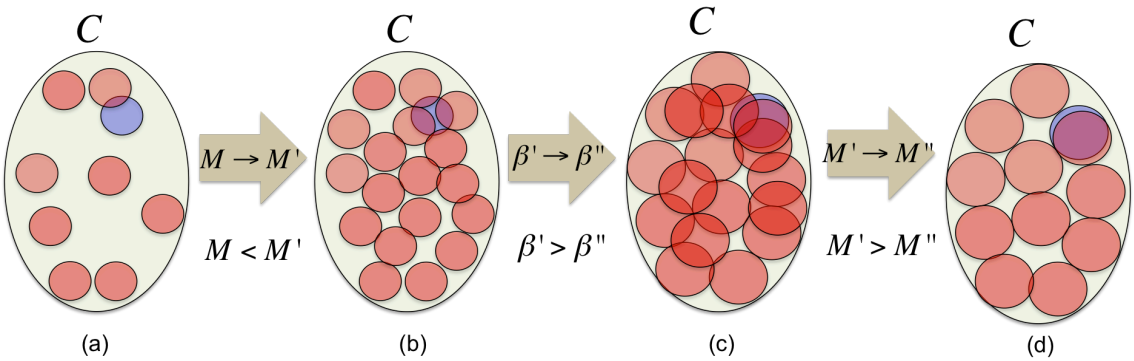

Figure 1.12: Learning protocol: maximizing learning rate

We have arrived at a fundamental trade-off between, on the one hand, the expressiveness of our code (the cardinality of $\mathcal{T}_M^h$) and the resolution of our decoding mechanism (the parameter β) and on the other, the stability of the learning protocol measured by the probability of error, itself a function of the *quality* of the overlap $\max_{j\in\{1,\dots,M\}} \Delta Z_\beta^j$.

Suppose the sender sends the hypothesis $c_j = t_j^h(c_D)$ corresponding to the dataset $t_j^D(X')$.

We want to compare the approximation sets $w_\beta(c, t_j^D(X'))$ and $w_\beta(c, t_j^D(X''))$ to get an idea of what needs to happen to be able to achieve the maximum learning rate. For a given resolution β:

1 We want to be able to increase the size of the message set, taking care that the approximation sets $w_\beta(c, t_j^D(X'))$ don't become cluttered. We can achieve this by maximizing:

$$\frac{|\mathcal{C}|}{Z_\beta(t_j^D(X'))} \tag{1.6.0.29}$$

This criterion will tend to make $Z_\beta(t_j^D(X'))$ small.

2 We want to make sure that the *quality* of the overlap $\Delta Z_\beta(t_j^D(X'), t_j^D(X''))$ is good. However we need a relative measure since low resolution communication will in general lead to a bigger $\Delta Z_\beta(t_j^D(X'), t_j^D(X''))$. We can achieve this by maximizing:

$$\frac{\Delta Z_\beta(t_j^D(X'), t_j^D(X''))}{Z_\beta(t_j^D(X''))} \tag{1.6.0.30}$$

This criterion will tend to make $Z_\beta(t_j^D(X''))$ small and $\Delta Z_\beta(t_j^D(X'), t_j^D(X''))$ large.

Now notice that for a given X', X'' and β these two quantities don't depend on the transformation t_j^D due to assumption 1.6.0.12 so that we may assume without loss of generality that $t_j^h(c) = c$ and $t_j^D(X) = X$. So for a given X' and X'' the bigger $\frac{|C|}{Z_\beta(X')}$ and $\frac{\Delta Z_\beta(X', X'')}{Z_\beta(X'')}$ the larger the learning rate will be. This gives us a qualitative notion of why generalization capacity defines the maximum learning rate of our learning algorithm.

We now analyze the probability of error of the above Learning Protocol to more formally understand the role of generalization capacity.

Recall what the random quantities in the Learning Protocol are:

- X',$X'' \sim p(x|c)$
- $\mathcal{T}_M^h = \{t_1^h, ..., t_M^h\}$, where $t_j^h \sim p(t^h) = \frac{1}{|\mathcal{C}|}$
- $\tilde{t}_s^h \sim p_s(t^h) = \frac{1}{M}$

Now let $J_{s^-} := \{1, ..., M\} \setminus \{s\}$. We have that:

$$\begin{aligned}
\mathbb{P}(\hat{t}^h \neq \tilde{t}_s^h | \tilde{t}_s^h) &= \mathbb{P}(\max_{j \in J_{s^-}} \Delta Z_\beta^j \geq \Delta Z_\beta^s | \tilde{t}_s^h) && (1.6.0.31)\\
&= \mathbb{P}(OR_{j \in J_{s^-}} (\Delta Z_\beta^j \geq \Delta Z_\beta^s) | \tilde{t}_s^h) && (1.6.0.32)\\
&\leq \sum_{j \in J_{s^-}} \mathbb{P}(\Delta Z_\beta^j \geq \Delta Z_\beta^s | \tilde{t}_s^h) && (1.6.0.33)\\
&= \sum_{j \in J_{s^-}} \mathbb{P}\left(\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s} \geq 1 | \tilde{t}_s^h\right) && (1.6.0.34)
\end{aligned}$$

Where we have used the union bound to establish the inequality. Note that $\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s}$ is a random variable depending on X', X'', $\mathcal{T}_M^h$ and $\widetilde{t}_s^h$. Using the Markov inequality and the independence of (X', X'') and t_j^h we see the following:

$$\mathbb{P}(\hat{t}^h \neq \widetilde{t}_s^h | \widetilde{t}_s^h) \leq \sum_{j \in J_{s^-}} \mathbb{P}\Big(\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s} \geq 1 | \widetilde{t}_s^h\Big) \tag{1.6.0.35}$$

$$\overset{(1)}{\leq} \sum_{j \in J_{s^-}} \mathbb{E}_{(X', X'', \mathcal{T}_M^h, \widetilde{t}_s^h)} \Big[\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s} | \widetilde{t}_s^h\Big] \tag{1.6.0.36}$$

$$\overset{(2)}{=} \sum_{j \in J_{s^-}} \mathbb{E}_{(X', X'', \mathcal{T}_M^h)} \Big[\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s}\Big] \tag{1.6.0.37}$$

$$\overset{(3)}{=} \sum_{j \in J_{s^-}} \mathbb{E}_{(X', X'')} \Big\{ \mathbb{E}_{\mathcal{T}_M^h} \Big[\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s}\Big] | \mathcal{T}_M^h \Big\} \tag{1.6.0.38}$$

$$\overset{(4)}{=} \sum_{j \in J_{s^-}} \mathbb{E}_{(X', X'')} \Big\{ \mathbb{E}_{t_j^h} \Big[\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s}\Big] | t_j^h \Big\} \tag{1.6.0.39}$$

$$\overset{(5)}{=} \sum_{j \in J_{s^-}} \mathbb{E}_{(X', X'')} \Big\{ \mathbb{E}_{t_j^h} \Big[\frac{\Delta Z_\beta^j}{\Delta Z_\beta^s}\Big] \Big\} \tag{1.6.0.40}$$

$$\overset{(6)}{=} \sum_{j \in J_{s^-}} \mathbb{E}_{(X', X'')} \Big\{ \frac{1}{\Delta Z_\beta^s} \mathbb{E}_{t_j^h} \Big[\Delta Z_\beta^j\Big] \Big\} \tag{1.6.0.41}$$

$$\overset{(7)}{=} \sum_{j \in J_{s^-}} \mathbb{E}_{(X', X'')} \Big\{ \frac{1}{\Delta Z_\beta^s} \mathbb{E}_{t_j^h} \Big[\Delta Z_\beta^j | X', X''\Big] \Big\} \tag{1.6.0.42}$$

$$\overset{(8)}{=} \mathbb{E}_{(X', X'')} \Big\{ \frac{1}{\Delta Z_\beta^s} \sum_{j \in J_{s^-}} \mathbb{E}_{t_j^h} \Big[\Delta Z_\beta^j | X', X''\Big] \Big\} \tag{1.6.0.43}$$

$$\overset{(9)}{=} (M-1) \mathbb{E}_{(X', X'')} \Big\{ \frac{1}{\Delta Z_\beta^s} \mathbb{E}_{t_j^h} \Big[\Delta Z_\beta^j | X', X''\Big] \Big\} \tag{1.6.0.44}$$

Where:

- (1) is due to the Markov inequality,
- (2) is due to the fact that $\widetilde{t}_s^h$ is given so there is no need to integrate over it,
- (3) is due to the chain rule of probability,
- (4) is due to the fact that ΔZ_β^j is independent of t_i^h for $i \neq j$,
- (5) and (7) are due to the fact that (X', X'') is independent of t_j^h,
- (6) is due to the fact that ΔZ_β^s is independent of t_j^h for $j \neq s$,
- (8) is due to the linearity of expectations, and
- (9) is due to the fact that t_j^h are identically distributed for $j \in \{1, ..., M\}$.

Now

$$\mathbb{E}_{t_j^h}\left[\Delta Z_\beta^j | X', X''\right] = \sum_{j=1}^{|\mathcal{T}_M^h|} \frac{1}{|\mathcal{T}_M^h|} \Delta Z_\beta^j \tag{1.6.0.45}$$

$$= \frac{1}{|\mathcal{T}_M^h|} \sum_{j=1}^{|\mathcal{T}_M^h|} \sum_{c \in \mathcal{C}} w_\beta(c, t_j^D(X')) w_\beta(c, \tilde{t}_s^D(X'')) \tag{1.6.0.46}$$

$$= \frac{1}{|\mathcal{T}_M^h|} \sum_{c \in \mathcal{C}} w_\beta(c, \tilde{t}_s^D(X'')) \sum_{j=1}^{|\mathcal{T}_M^h|} w_\beta(c, t_j^D(X')) \tag{1.6.0.47}$$

$$= \frac{1}{|\mathcal{T}_M^h|} \sum_{c \in \mathcal{C}} w_\beta((\tilde{t}_s^h)^{-1}(c), X'') \sum_{j=1}^{|\mathcal{T}_M^h|} w_\beta((t_j^h)^{-1}(c), X') \tag{1.6.0.48}$$

Where for the last equivalence we have used assumption 1.6.0.12. Using the properties of 1.6.0.10 we can establish the following:

$$\mathbb{E}_{t_j^h}\left[\Delta Z_\beta^j | X', X''\right] = \frac{1}{|\mathcal{T}_M^h|} \sum_{c \in \mathcal{C}} w_\beta((\tilde{t}_s^h)^{-1}(c), X'') \sum_{j=1}^{|\mathcal{T}_M^h|} w_\beta((t_j^h)^{-1}(c), X') \tag{1.6.0.49}$$

$$= \frac{1}{|\mathcal{T}_M^h|} \sum_{c \in \mathcal{C}} w_\beta((\tilde{t}_s^h)^{-1}(c), X'') \sum_{c \in \mathcal{C}} w_\beta(c, X') \tag{1.6.0.50}$$

$$= \frac{Z_\beta(X')}{|\mathcal{T}_M^h|} \sum_{c \in \mathcal{C}} w_\beta((\tilde{t}_s^h)^{-1}(c), X'') \tag{1.6.0.51}$$

$$= \frac{Z_\beta(X')}{|\mathcal{T}_M^h|} \sum_{c \in \mathcal{C}} w_\beta(c, X'') \tag{1.6.0.52}$$

$$= \frac{Z_\beta(X') Z_\beta(X'')}{|\mathcal{T}_M^h|} \tag{1.6.0.53}$$

So we have that

$$\mathbb{P}(\hat{t}^h \neq \tilde{t}_s^h | \tilde{t}_s^h) \leq (M-1) \mathbb{E}_{(X', X'')} \left\{ \frac{1}{\Delta Z_\beta^s} \mathbb{E}_{t_j^h}\left[\Delta Z_\beta^j | X', X''\right] \right\} \tag{1.6.0.54}$$

$$= (M-1) \mathbb{E}_{(X', X'')} \left\{ \frac{Z_\beta(X') Z_\beta(X'')}{|\mathcal{T}_M^h| \Delta Z_\beta^s} \right\} \tag{1.6.0.55}$$

Again using the properties of 1.6.0.10 we have that:

$$\Delta Z^s_\beta = \sum_{c\in\mathcal{C}} w_\beta(c, \widetilde{t}^{\mathcal{D}}_s(X'))w_\beta(c, \widetilde{t}^{\mathcal{D}}_s(X'')) \tag{1.6.0.56}$$

$$= \sum_{c\in\mathcal{C}} w_\beta((\widetilde{t}^h_s)^{-1}(c), X')w_\beta((\widetilde{t}^h_s)^{-1}(c), X'') \tag{1.6.0.57}$$

$$= \sum_{c\in\mathcal{C}} w_\beta(c, X')w_\beta(c, X'') \tag{1.6.0.58}$$

$$= \Delta Z_\beta(X', X'') \tag{1.6.0.59}$$

with which

$$\mathbb{P}(\hat{t}^h \neq \widetilde{t}^h_s|\widetilde{t}^h_s) \leq (M-1)\mathbb{E}_{(X',X'')}\left\{\frac{Z_\beta(X')Z_\beta(X'')}{\Delta Z^s_\beta}\right\} \tag{1.6.0.60}$$

$$= (M-1)\mathbb{E}_{(X',X'')}\left\{\frac{Z_\beta(X')Z_\beta(X'')}{|\mathcal{T}^h_M|\Delta Z_\beta(X', X'')}\right\} \tag{1.6.0.61}$$

$$= (M-1)\mathbb{E}_{(X',X'')}\left\{\frac{Z_\beta(X')Z_\beta(X'')}{|\mathcal{C}|\Delta Z_\beta(X', X'')}\right\} \tag{1.6.0.62}$$

$$= \mathbb{E}_{(X',X'')}[e^{\log(M-1)-I_\beta}] \tag{1.6.0.63}$$

So error free learning is possible as long as, *on average*, for data sets X' and X'' it holds that $\log(M-1) < I_\beta$.

Finally we mention that since the generalization capacity defines the maximum learning rate of a cost function $\hat{R}(c, X)$ it can be used as a criterion for deciding which cost function to use: simply use the cost function with highest generalization capacity. We will see an application of this in Section 3.6.

Chapter 2

Mean localization

2.1 Generalization capacity

Recall from Section 1.4 that the empirical risk function $\hat{R}(\mu, X)$ is defined with respect to a loss function $\rho(\mu, X)$. We will mostly deal with the *square loss* function:

$$\rho_\mu(x) = ||x - \mu||_2^2 = \sum_{j=1}^{d}(x_j - \mu)^2 \tag{2.1.0.1}$$

We calculate the empirical risk, which is our cost function in the context of ASC, with respect to this loss function.

$$\hat{R}(\mu, X) = \frac{1}{n}\sum_{i=1}^{n}\rho_\mu(x_i) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{d}(x_{ij} - \mu_j)^2 = \sum_{j=1}^{d}\frac{1}{n}\sum_{i=1}^{n}x_{ij}^2 - 2\mu_j x_{ij} + \mu_j^2 \tag{2.1.0.2}$$

$$= \sum_{j=1}^{d}\left\{\frac{1}{n}\sum_{i=1}^{n}x_{ij}^2 - \frac{1}{n}\sum_{i=1}^{n}2\mu_j x_{ij} + \frac{1}{n}\sum_{i=1}^{n}\mu_j^2\right\} \tag{2.1.0.3}$$

$$= \sum_{j=1}^{d}\left\{\frac{1}{n}\sum_{i=1}^{n}x_{ij}^2 - 2\mu_j\bar{X}_j + \mu_j^2\right\} \tag{2.1.0.4}$$

$$\overset{\mu}{\propto} \sum_{j=1}^{d}\left\{\mu_j^2 - 2\mu_j\bar{X}_j\right\} \tag{2.1.0.5}$$

$$\overset{\mu}{\propto} \sum_{j=1}^{d}\left\{\mu_j^2 - 2\mu_j\bar{X}_j + \bar{X}_j^2\right\} \tag{2.1.0.6}$$

$$= \sum_{j=1}^{d}(\mu_j - \bar{X}_j)^2 = ||\mu - \bar{X}||_2^2 \tag{2.1.0.7}$$

With which we can define the cost function as:

$$\hat{R}(\mu, X) = ||\mu - \bar{X}||_2^2 \tag{2.1.0.8}$$

Recall that in the problem at hand $\mathcal{C} = \mathbb{B}^d$ so that:

$$I = \max_{\beta \in \mathbb{R}^+} \mathbb{E}_{(X', X'')} I_\beta = \max_{\beta \in \mathbb{R}^+} \mathbb{E}_{(X', X'')} \log \frac{|\mathbb{B}^d| \Delta Z_\beta(X', X'')}{Z_\beta(X') Z_\beta(X'')} \tag{2.1.0.9}$$

$$= \log 2^d + \mathbb{E}_{(X', X'')} \log \Delta Z_\beta(X', X'') - \mathbb{E}_{X'} \log Z_\beta(X') - \mathbb{E}_{X''} \log Z_\beta(X'') \tag{2.1.0.10}$$

Where:

- $Z_\beta(X) = \sum_{\mu \in \mathbb{B}^d} e^{-\beta \hat{R}(\mu, X)} = \sum_{\mu \in \mathbb{B}^d} e^{-\beta ||\mu - \bar{X}||_2^2}$ and
- $\Delta Z_\beta(X', X'') = \sum_{\mu \in \mathbb{B}^d} e^{-\beta \{\hat{R}(\mu, X') + \hat{R}(\mu, X'')\}} = \sum_{\mu \in \mathbb{B}^d} e^{-\beta \{||\mu - \bar{X}_1||_2^2 + ||\mu - \bar{X}_2||_2^2\}}$

Since we don't know the distribution of $Z_\beta(X)$ or $\Delta Z_\beta(X', X'')$ we will use simulation to estimate I. Before describing the simulation algorithm notice that I is a function of the two data set means $\bar{X}_1$ and $\bar{X}_2$ and since X has a normal distribution:

$$X \sim N\Big(\mu^0, \sigma^2 I_d\Big) \Rightarrow \bar{X} \sim N\Big(\mu^0, \frac{\sigma^2}{n} I_d\Big) \tag{2.1.0.11}$$

If we let $\xi = \frac{\bar{X} - \mu^0}{\sigma / \sqrt{n}} \sim N(0, I_d)$ then $\bar{X} = \mu^0 + \frac{\sigma}{\sqrt{n}} \xi$ and we have that:

$$I = \max_{\beta \in \mathbb{R}^+} \mathbb{E}_{(\xi_1, \xi_2)} \log \frac{|\mathbb{B}^d| \Delta Z_\beta(\xi_1, \xi_2)}{Z_\beta(\xi_1) Z_\beta(\xi_2)} \tag{2.1.0.12}$$

Where:

- $\hat{R}(\mu, \xi) = ||\mu - \mu^0 - \frac{\sigma}{\sqrt{n}} \xi||_2^2$,
- $w_\beta(\mu, \xi) = e^{-\beta \hat{R}(\mu, \xi)}$
- $P_G(\mu; \beta, \xi) = \frac{w_\beta(\mu, \xi)}{\sum_{\mu \in \mathbb{B}^d} w_\beta(\mu, \xi)}$
- $Z_\beta(\xi) = \sum_{\mu \in \mathbb{B}^d} e^{-\beta ||\mu - \mu^0 - \frac{\sigma}{\sqrt{n}} \xi||_2^2}$ and
- $\Delta Z_\beta(\xi_1, \xi_2) = \sum_{\mu \in \mathbb{B}^d} e^{-\beta \{||\mu - \mu^0 - \frac{\sigma}{\sqrt{n}} \xi_1||_2^2 + ||\mu - \mu^0 - \frac{\sigma}{\sqrt{n}} \xi_2||_2^2\}}$

Using this alternate expression we can reduce the number of simulations by a factor of n. Notice that varying both n and σ doesn't make sense since $\frac{\sigma^2}{n}$ characterizes the variance of the data set. For this reason we leave $n = 100$ fixed and only vary σ in our simulation experiments.

2.2 Exhaustive algorithm

To estimate the generalization capacity for the mean localization problem, for a given μ^0 and σ, take the following steps:

1. Choose a grid of relevant β values: $\underline{\beta} = (\beta_1, ..., \beta_l)$
2. For $i = 1$ to m
   a. Simulate $\xi_1^i, \xi_2^i \sim N(0, I_d)$
   b. For $k = 1$ to l
      - Calculate information content:

$$\hat{I}^i_{\beta_k} = \log \frac{2^d \Delta Z_{\beta_k}(\xi_1^i, \xi_2^i)}{Z_{\beta_k}(\xi_1^i) Z_{\beta_k}(\xi_2^i)} \tag{2.2.0.1}$$

3. For $k = 1$ to l
   - Estimate mean information content:

$$\bar{I}_{\beta_k} = \frac{1}{m} \sum_{i=1}^{m} \hat{I}^i_{\beta_k} \tag{2.2.0.2}$$

   - Estimate generalization capacity:

$$\hat{I} = \max_{k \in \{1,...,l\}} \bar{I}_{\beta_k} \tag{2.2.0.3}$$

Remark. When dealing with real data we don't know what μ^0 is. This doesn't matter since GC is independent of any particular hypothesis, rather it depends on the data generating mechanism and on the cost function $\hat{R}(\mu, X)$. Provided we can simulate from the model that we assume generated the data we will always be able to estimate GC by simulating from the model for an arbitrary set of parameters μ^0. By calculating GC we obtain a β^* that resolves the resolution-stability trade off. We may then use β^* on the real data to obtain an appropriate approximation set of hypotheses. Alternatively we may compare the GC associated to different cost functions and choose the cost function with the highest GC.

2.3 Simulation results

We used the following parameters for the simulation experiments:

- $d = 8$,
- $m = 200$,
- 100 different β values from 0.01 to 20, and
- 30 different noise levels σ from 0.1 to 10.

To check simulation results made sense we first calculated the Gibbs distribution $P_G(c;\beta^*,\xi)$ distribution over $\mathcal{C}=\mathbb{B}^d$, where β^* is the resolution parameter that allows generalization capacity to be reached. Since $d=8$ $P_G(c;\beta^*,\xi)$ corresponds to a vector with $2^8=258$ entries. To display the results in an easy to read fashion we aggregated this vector to produce a component-wise Gibbs distribution:

$$\mathbb{P}_G(\mu_j^0=1|\beta^*)=\sum_{\mu\in\mathbb{B}^d}P_G(\mu;\beta^*,\xi)\mathbb{1}_{\{\mu_j=1\}} \tag{2.3.0.1}$$

The following graph is an illustration of the $\mathbb{P}_G(\mu_j^0=1|\beta^*)$ estimate for different noise levels.

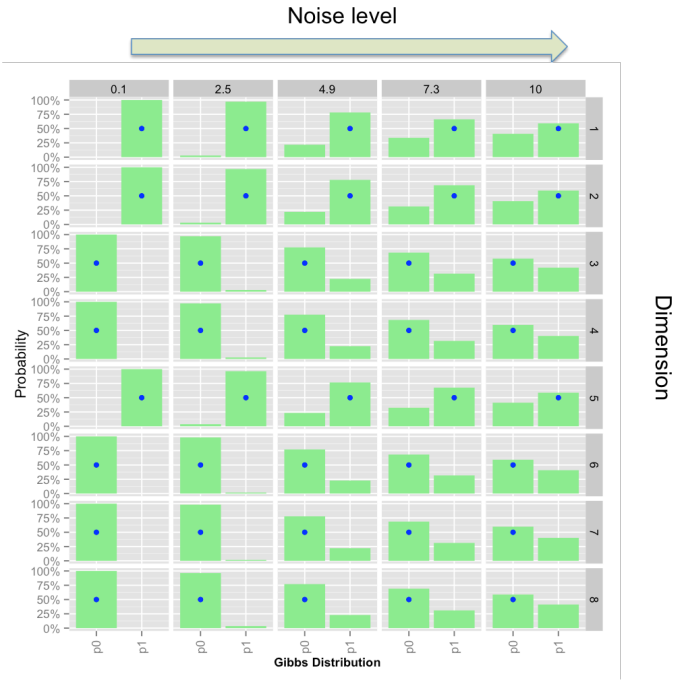

Figure 2.1: Component-wise Gibbs distribution

The blue dots show the true value of μ^0 for each of its components. The lower the noise σ the less uncertainty about the value of the μ^0 we have. We next show the average information content for different resolutions β and noise levels σ.

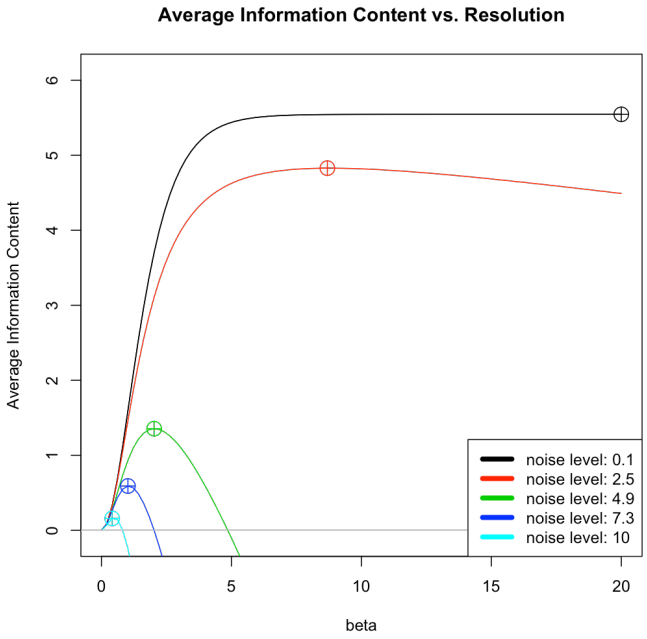

Figure 2.2: Average information content

The crossed circles represent the pairs (β^*, I) where average information content is maximized. The lower the noise level the higher the generalization capacity is. For *low* noise levels, sucha as $\sigma = 0.1$, we can obtain gains in average information content the higher the resolution β (albeit at a diminishing rate) i.e. generalization capacity is basically an increasing function of resolution for these noise levels. This means that for such low noise we can let $\beta \to \infty$ and obtain the empirical risk minimizer. For *medium* range noise levels, such as $\sigma = 4.9$, once we go past the resolution threshold β^*, the average information content decreases dramatically i.e. generalization capacity is a convex function of resolution for these noise levels. Here we see the resolution-stability trade-off clearly. For this level of noise we can only decrease our approximation sets to a certain size parametrized by β^* before we start to get very unstable sets with little information.

We now show the generalization capacity for different noise levels.

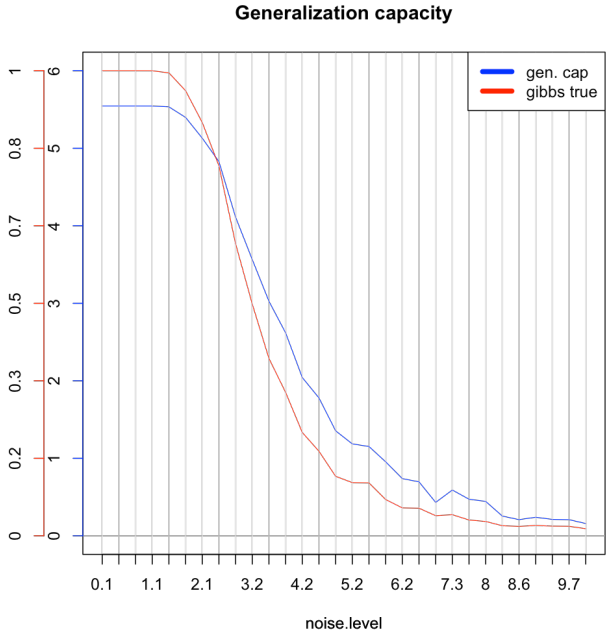

Figure 2.3: Generalization Capacity

The blue line shows the generalization capacity while the red line shows the *true* Gibbs probabilty, i.e. $P_G(\mu^0; \beta^*, \xi)$. As expected the generalization capacity decreases toward zero as the noise level becomes so big as to completely drown out the signal μ^0. Notice that for noise levels $\sigma < 1.3$, $P_G(\mu^0; \beta^*, \xi) = 1$ and $I = d \log 2 = 8 \log 2 \approx 5.55$ indicating that we can completely recover the signal μ^0 with the empirical risk minimizer ($I = d \log 2 \Rightarrow Z_\beta(\xi_1) Z_\beta(\xi_2) = \Delta Z_\beta(\xi_1, \xi_2)$!).

2.4 Log-sum-exp trick

We can express the information content as

$$I_\beta = \log |\mathcal{C}| + \log \Delta Z_\beta(\xi_1, \xi_2) - \log Z_\beta(\xi_1) - \log Z_\beta(\xi_2). \tag{2.4.0.1}$$

In turn, we can express the log partition functions as:

$$\log Z_\beta(\xi) = \log \sum_{i=1}^{\mathcal{C}} e^{a_i} \tag{2.4.0.2}$$

$$\log \Delta Z_\beta(\xi_1, \xi_2) = \log \sum_{i=1}^{\mathcal{C}} e^{b_i} \tag{2.4.0.3}$$

where:

- $a_i = -\beta \hat{R}(\xi, \mu_i) < 0$,

- $b_i = -\beta(\hat{R}(\xi_1, \mu_i) + \hat{R}(\xi_2, \mu_i)) < 0$ and
- $\mathcal{C} = \{\mu_1, ..., \mu_{|\mathcal{C}|}\}$

When implementing the algorithm of Section 2.2 we may run into the problem that for very high resoluton values β, $\sum_{i=1}^{\mathcal{C}} e^{a_i}$ and $\sum_{i=1}^{\mathcal{C}} e^{b_i}$ are so small that they are represented as 0 using limited-precision, floating point numbers. This means $\log Z_\beta(\xi)$ and $\log \Delta Z_\beta(\xi_1, \xi_2)$ are represented as $-\infty$ with which our calculation of I_β breaks down. To solve this *underflow* problem we use the log-sum-exp trick. Let:

$$A = \max_{i \in \{1,...,|\mathcal{C}|\}} a_i \tag{2.4.0.4}$$

Then,

$$\log Z_\beta(\xi) = \log \sum_{i=1}^{\mathcal{C}} e^{a_i} = \log e^A \sum_{i=1}^{\mathcal{C}} e^{a_i - A} = A + \log \sum_{i=1}^{\mathcal{C}} e^{a_i - A} \tag{2.4.0.5}$$

If we calculate $\log Z_\beta(\xi)$ using the last expression, this solves the underflow problem since $\sum_{i=1}^{\mathcal{C}} e^{a_i - A} > 1$ so that $\log \sum_{i=1}^{\mathcal{C}} e^{a_i - A} > 0$.

2.5 Variance reduction with common random numbers

If we want to estimate by simulation:

$$\mu_{f-g} = \mathbb{E}_X[f(X) - g(X)] = E_X[f(X)] - E_X[g(X)] \tag{2.5.0.1}$$

We can either:

a. Generate $2m$ i.i.d. realizations of X: $X_1 = \{x_{11}, ..., x_{1m}\}$,$X_2 = \{x_{21}, ..., x_{2m}\}$. Then:

$$\hat{\mu}_f^1 = \hat{\mathbb{E}}_X^1[f(X)] = \frac{1}{m} \sum_{i=1}^{m} f(x_{1i}) \tag{2.5.0.2}$$

$$\hat{\mu}_g^2 = \hat{\mathbb{E}}_X^2[g(X)] = \frac{1}{m} \sum_{i=1}^{m} g(x_{2i}) \tag{2.5.0.3}$$

$$\hat{\mathbb{E}}_X^{12}[f(X) - g(X)] = \hat{\mathbb{E}}_X^1[f(X)] - \hat{\mathbb{E}}_X^2[g(X)] = \hat{\mu}_f^1 - \hat{\mu}_g^2 \tag{2.5.0.4}$$

or we can,

b. Generate m i.i.d. realizations of X: $X_1 = \{x_{11}, ..., x_{1m}\}$. Then:

$$\hat{\mu}_f = \hat{\mathbb{E}}_X[f(X)] = \frac{1}{m} \sum_{i=1}^{m} f(x_{1i}) \tag{2.5.0.5}$$

$$\hat{\mu}_g = \hat{\mathbb{E}}_X[g(X)] = \frac{1}{m} \sum_{i=1}^{m} g(x_{1i}) \tag{2.5.0.6}$$

$$\hat{\mathbb{E}}_X[f(X) - g(X)] = \hat{\mathbb{E}}_X[f(X)] - \hat{\mathbb{E}}_X[g(X)] = \hat{\mu}_f - \hat{\mu}_g \tag{2.5.0.7}$$

The second method is an example of *Common Random Numbers* (CRNs) since we use the same pseudo-random numbers to estimate $\mathbb{E}_X[f(X)]$ and $\mathbb{E}_X[g(X)]$.

CRNs is a *variance reduction* technique, although, strictly speaking, it does not always succeed in reducing the variance of an estimator. To see when it might succeed we compare the variance of $\hat{\mu}_f^1 - \hat{\mu}_g^2$ and $\hat{\mu}_f - \hat{\mu}_g$:

$$\mathbb{V}_X[\hat{\mu}_f^1 - \hat{\mu}_g^2] = \mathbb{V}_X\left[\frac{1}{m}\sum_{i=1}^{m} f(x_{1i}) - \frac{1}{m}\sum_{i=1}^{m} g(x_{2i})\right] \tag{2.5.0.8}$$

$$= \mathbb{V}_X\left[\frac{1}{m}\sum_{i=1}^{m} (f(x_{1i}) - g(x_{2i}))\right] \tag{2.5.0.9}$$

$$= \frac{1}{m^2}\sum_{i=1}^{m} \mathbb{V}_X(f(x_{1i}) - g(x_{2i})) \tag{2.5.0.10}$$

$$= \frac{1}{m^2}\sum_{i=1}^{m} (\mathbb{V}_X[f(X)] + \mathbb{V}_X[g(X)]) \tag{2.5.0.11}$$

$$= \frac{1}{m}(\mathbb{V}_X[f(X)] + \mathbb{V}_X[g(X)]) \tag{2.5.0.12}$$

$$\mathbb{V}_X[\hat{\mu}_f - \hat{\mu}_g] = \mathbb{V}_X\left[\frac{1}{m}\sum_{i=1}^{m} f(x_{1i}) - \frac{1}{m}\sum_{i=1}^{m} g(x_{1i})\right] \tag{2.5.0.13}$$

$$= \mathbb{V}_X\left[\frac{1}{m}\sum_{i=1}^{m} (f(x_{1i}) - g(x_{1i}))\right] \tag{2.5.0.14}$$

$$= \frac{1}{m^2}\sum_{i=1}^{m} \mathbb{V}_X[f(x_{1i}) - g(x_{1i})] \tag{2.5.0.15}$$

$$= \frac{1}{m}\mathbb{V}_X[f(X) - g(X)] \tag{2.5.0.16}$$

$$= \frac{1}{m}\mathbb{V}_X[f(X)] + \mathbb{V}_X[g(X)] - 2Cov(f(X), g(X)) \tag{2.5.0.17}$$

$$\tag{2.5.0.18}$$

If f and g are both either monotonically non-decreasing or mononotonically non-increasing then:

$$Cov(X, X) = \mathbb{V}_X[X] > 0 \Rightarrow Cov(f(X), g(X)) \geq 0 \tag{2.5.0.19}$$

$$\Rightarrow \mathbb{V}_X[\hat{\mu}_f - \hat{\mu}_g] \leq \mathbb{V}_X[\hat{\mu}_f^1 - \hat{\mu}_g^2] \tag{2.5.0.20}$$

We now obtain a different expression for $Z_\beta(\xi)$ and $\Delta Z_\beta(\xi_1, \xi_2)$ in order to use the CRN technique for the estimation of $\mathbb{E}_{(\xi_1,\xi_2)}[I_\beta]$. We start with an alternative, but equivalent (proportional in μ) cost function $\hat{R}(\mu, X)$. Using that $\mu_j \in \{0, 1\}$ we have that:

$$\hat{R}(\mu, X) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{d}(x_{ij}-\mu_j)^2 = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{d}\mu_j(x_{ij}-1)^2 + (1-\mu_j)x_{ij}^2 \tag{2.5.0.21}$$

$$= \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{d}\mu_j(1-2x_{ij}) + x_{ij}^2 \tag{2.5.0.22}$$

$$\overset{\mu}{\propto} \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{d}\mu_j(1-2x_{ij}) \tag{2.5.0.23}$$

$$= \sum_{j=1}^{d}\mu_j\frac{1}{n}\sum_{i=1}^{n}(1-2x_{ij}) \tag{2.5.0.24}$$

$$= \sum_{j=1}^{d}\mu_j(1-2\bar{x}_j) = \mu^T(1-2\bar{x}) \tag{2.5.0.25}$$

So that $\mu^T(1-2\bar{x}) \overset{\mu}{\propto} ||\mu-\bar{x}||_2^2$.

Again, using that $\bar{X} = \frac{\sigma}{\sqrt{n}}\xi + \mu^0$ we have that:

$$\hat{R}(\mu, \xi) = \mu^T(1 - 2\frac{\sigma}{\sqrt{n}}\xi - 2\mu^0) \tag{2.5.0.26}$$

$$\hat{R}(\mu, \xi_1) + \hat{R}(\mu, \xi_2) = 2\mu^T(1 - \frac{\sigma}{\sqrt{n}}(\xi_1+\xi_2) - 2\mu^0) \tag{2.5.0.27}$$

Since $\frac{1}{\sqrt{2}}(\xi_1+\xi_2) \sim N(0, I_d)$ we have that:

$$Z_\beta(\xi) = \sum_{\mu\in\mathbb{B}^d} e^{-\beta\hat{R}(\mu,\xi)} \tag{2.5.0.28}$$

$$\Delta Z_\beta(\xi_1, \xi_2) = \Delta Z_\beta(\xi) = \sum_{\mu\in\mathbb{B}^d} e^{-\beta\hat{R}'(\mu,\xi)} \tag{2.5.0.29}$$

Where:

- $\hat{R}(\mu, \xi) := \mu^T(1 - 2\frac{\sigma}{\sqrt{n}}\xi - 2\mu^0)$
- $\hat{R}'(\mu, \xi) := 2\mu^T(1 - \frac{\sigma}{\sqrt{n}}\sqrt{2}\xi - 2\mu^0)$

Using these expressions we implemented 3 different CRN algorithms to estimate I_β and I, and compared the variance estimate $\hat{\mathbb{V}}_{(\xi_1,\xi_2)}[\hat{I}_{\beta^*}]$ of each, where $\beta^* = \underset{\beta\in\{\beta_1,\ldots,\beta_l\}}{\arg\max}\ \hat{\mathbb{E}}_{(\xi_1,\xi_2)}[\hat{I}_\beta]$.

All three algorithms generate $2m$ realizations of $\xi \sim N(0, I_d)$ but differ in wether they use them for the estimation of $\mathbb{E}_\xi[\log Z_\beta(\xi)]$, $\mathbb{E}_{(\xi_1,\xi_2)}[\log \Delta Z_\beta(\xi_1, \xi_2)] = \mathbb{E}_\xi[\log \Delta Z_\beta(\xi)]$ or both:

1. **CRN-1** Generate $2m$ realizations of $\xi \sim N(0, I_d)$. Use first m to calculate $\hat{\mathbb{E}}_{\xi_1}[\log Z_\beta(\xi_1)]$ and $\hat{\mathbb{E}}_{\xi_2}[\log Z_\beta(\xi_2)]$ and second m to calculate $\hat{\mathbb{E}}_\xi[\log \Delta Z_\beta(\xi)]$.

2. **CRN-2** Generate $2m$ realizations of $\xi \sim N(0, I_d)$. Use all $2m$ to calculate $\hat{\mathbb{E}}_{\xi_1}[\log Z_\beta(\xi_1)]$, $\hat{\mathbb{E}}_{\xi_2}[\log Z_\beta(\xi_2)]$ and $\hat{\mathbb{E}}_{(\xi_1,\xi_2)}[\log \Delta Z_\beta(\xi_1, \xi_2)]$.

3. **CRN-3** Generate $2m$ realizations of $\xi \sim N(0, I_d)$. Use first m to calculate $\hat{\mathbb{E}}_{\xi_1}[\log Z_\beta(\xi_1)]$, second m to calculate $\hat{\mathbb{E}}_{\xi_2}[\log Z_\beta(\xi_2)]$ and all $2m$ to calculate $\hat{\mathbb{E}}_{(\xi_1,\xi_2)}[\log \Delta Z_\beta(\xi_1, \xi_2)]$. This CRN algorithm is actually the algorithm proposed in Section 2.2. We simulate (ξ_1, ξ_2) m times and use it to estimate both $\mathbb{E}_{(\xi_1,\xi_2)}[\log Z_\beta(\xi_1) Z_\beta(\xi_2)]$ and $\mathbb{E}_{(\xi_1,\xi_2)}[\log \Delta Z_\beta(\xi_1, \xi_2)]$.

We used the following parameters for the simulation experiments:

- $d = 8$,
- $m = 100$,
- 100 different β values from 0.01 to 20, and
- 30 different noise levels σ from 0.1 to 4.

Figure 2.4b shows the estimation of the generalization capacity $\hat{I} = \max_{\beta \in \mathbb{R}^+} \hat{\mathbb{E}}_{(\xi_1,\xi_2)} \log \hat{I}_\beta$ and the square root of variance $\sqrt{\hat{\mathbb{V}}_{(\xi_1,\xi_2)}[\hat{I}_{\beta^*}]}$ using all 3 methods.

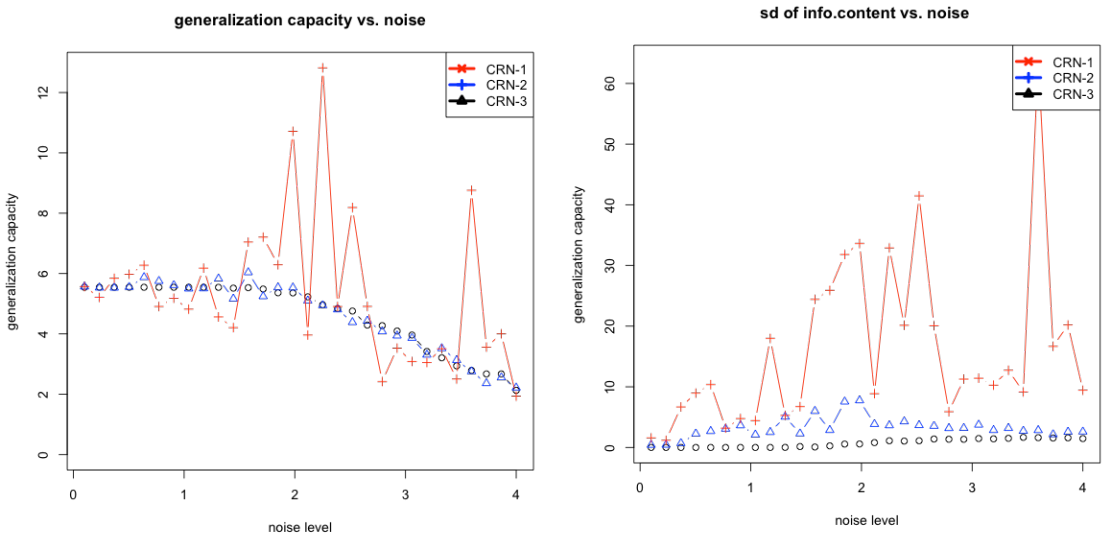

(a) $\hat{I}$ (b) $\sqrt{\hat{\mathbb{V}}_{(\xi_1,\xi_2)}[\hat{I}_{\beta^*}]}$

Figure 2.4: Calculation of generalization capacity with CRNs

Clearly method CRN-3, the method described in Section 2.2, has the least variance. This makes sense because if ξ_1, ξ_2, ξ_3 are i.i.d.:

a. $Cov(\log Z_\beta(\xi_1), \log Z_\beta(\xi_1)) > 0 \Rightarrow \mathbb{V}[\log Z_\beta(\xi_1) + \log Z_\beta(\xi_1)] > \mathbb{V}[\log Z_\beta(\xi_1) + \log Z_\beta(\xi_2)]$, so using the same random number to estimate $Z_\beta(\xi_1)$ and $Z_\beta(\xi_2)$ actually increases the variance, while,

b. $Cov(\log Z_\beta(\xi_1), \log \Delta Z_\beta(\xi_1, \xi_2)) > 0 \Rightarrow \mathbb{V}[\log \Delta Z_\beta(\xi_1, \xi_2) - \log Z_\beta(\xi_1)] < \mathbb{V}[\log \Delta Z_\beta(\xi_1, \xi_2) - \log Z_\beta(\xi_3)]$, so using the same random number to estimate $\Delta Z_\beta(\xi_1, \xi_2)$ and $Z_\beta(\xi_1)$ decreases the variance.

Chapter 3

Sparse mean localization

3.1 Exhaustive algorithm

For the sparse mean localization problem the hypothesis space is restricted to binary vectors with k entries equal to one:

$$\mathcal{C}^k := \{\mu \in \mathbb{B}^d : ||\mu||_1 = k\} \tag{3.1.0.1}$$

Notice that $|\mathcal{C}^k| = \binom{d}{k}$. In terms of the algorithm in Section 2.2, the only thing that changes is that the sums involved in $Z_\beta(\xi)$ and $\Delta Z_\beta(\xi_1, \xi_2)$ are over $\mathcal{C}^k$ instead of the entire $\mathcal{C} = \mathbb{B}^d$.

In the non-sparse case, to generate the vectors μ_i such that $\mathcal{C} = \{\mu_1, ..., \mu_{2^d}\}$ we can simply use the mapping $\kappa : \{0, ..., 2^d - 1\} \to \mathcal{C}$ where κ maps a positive integer $0 \leq i \leq 2^d - 1$ to its d length binary representation.

In the sparse case, to generate the vecors μ_i such that $\mathcal{C}^k = \{\mu_1, ..., \mu_{\binom{d}{k}}\}$ we use the mapping $\gamma : \{1, ..., \binom{d}{k}\} \to \mathcal{C}^k$ where γ maps a positive integer $1 \leq i \leq \binom{d}{k}$ to the i-th element of $\mathcal{C}^k$, assuming that $\mathcal{C}^k$ is in *lexicographical order*.

The algorithm shown below, adapted from **?** (pp. 27-29), can be used to produce the mapping γ. It is based on the fact that if $\mathcal{C}^k$ is in reverse lexicographical order then we can represent the number $1 \leq i \leq \binom{d}{k}$ as:

$$i = \binom{p_k}{k} + ... + \binom{p_2}{2} + \binom{p_1}{1} \tag{3.1.0.2}$$

Where $p_j \in \{1, ..., d\}$ gives the position of the j-th entry of μ_i that is equal to one. i.e. $\mu_{ip_j} = 1 \ \forall j \in \{1, ..., k\}$.

Algorithm to obtain μ_i

1. Initalize:
   - Set $\mu_{ij} \leftarrow 0 \ \forall j \in \{1, ..., d\}$.
   - Set $m \leftarrow i - 1$.
2. For j from k to 1:
   a. $p_j \leftarrow \underset{l \in \{t: \binom{t}{j} \leq m\}}{\arg\max} \binom{l}{j}$
   b. $m \leftarrow m - \binom{p_j}{j}$
   c. $\mu_{ip_j} \leftarrow 1$

3.2 Simulation results: exhaustive algorithm

We used the following parameters for the simulation experiments:

- $d = 10$,
- $k = 1, 2, 3, 4, 5$
- $m = 100$,
- 100 different β values from 0.01 to 20, and
- 30 different noise levels σ from 0.1 to 10.

First we show the information content for different values of k, σ and β. Since $d = 10$ we look at $k \in \{1, 2, ..., 5\}$ only. For $k \in \{6, 7, ..., 10\}$ the behavior will be equivalent and the only difference is that the role of 0 and 1 ($\mu_j \in \{0, 1\}$) is reversed.

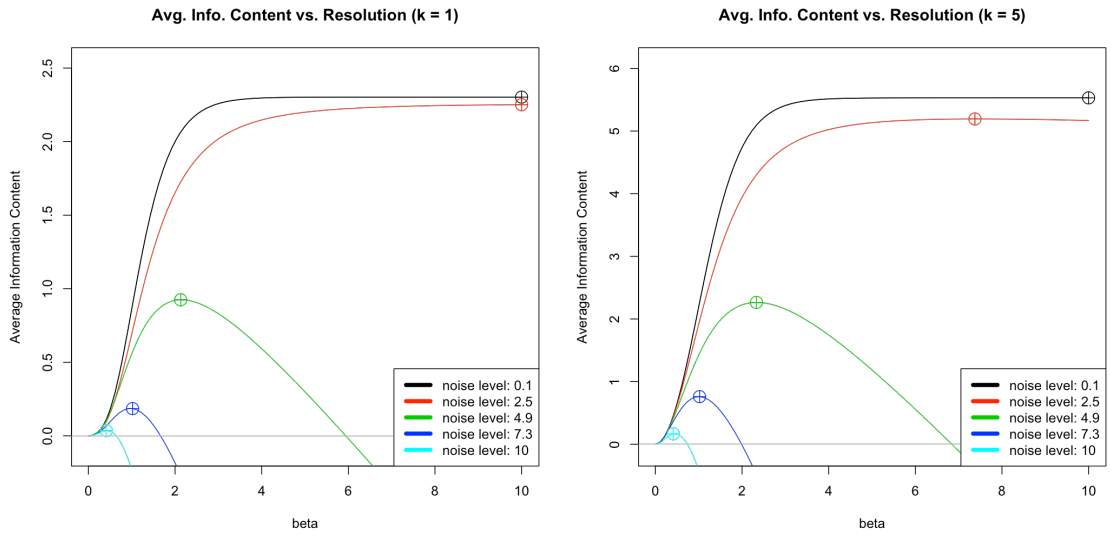

(a) $k = 1$ (b) $k = 5$

Figure 3.1: Average information content for different k

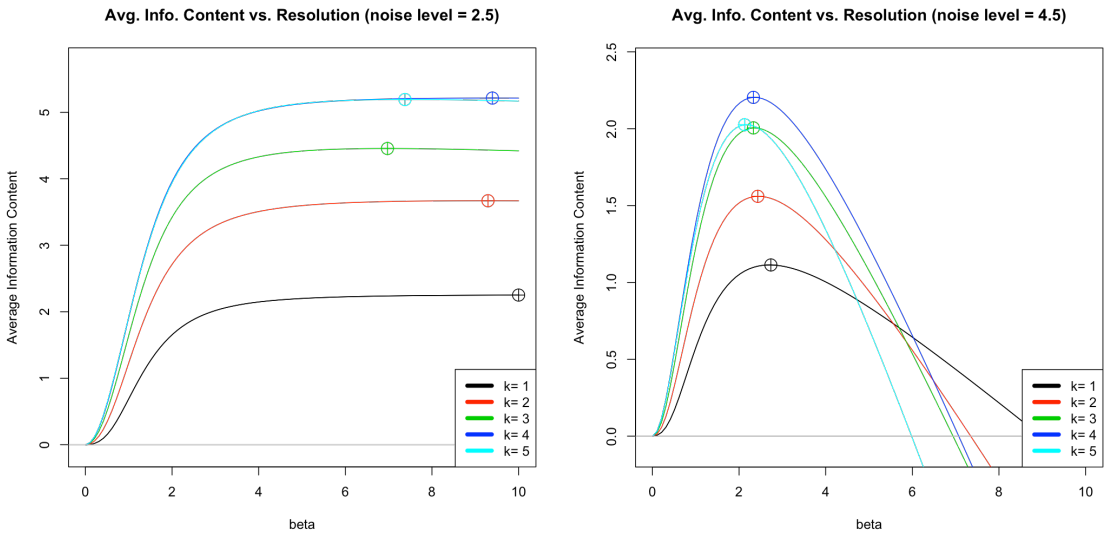

(a) $\sigma = 2.5$ (b) $\sigma = 4.5$

Figure 3.2: Average information content for different σ

For a given k we see that, as in the non-sparse case, for low noise levels, increases in resolution obtain diminishing gains in information content while for higher noise levels, once the optimum resolution β^* is surpassed, the information content decreases. Also notice that as k increases toward $\frac{d}{2} = 5$ the information content also increases. This is because the size of the hypothesis space is increasing in k from 0 to $\lceil \frac{d}{2} \rceil$. We now show the generalization capacity for different values of k.

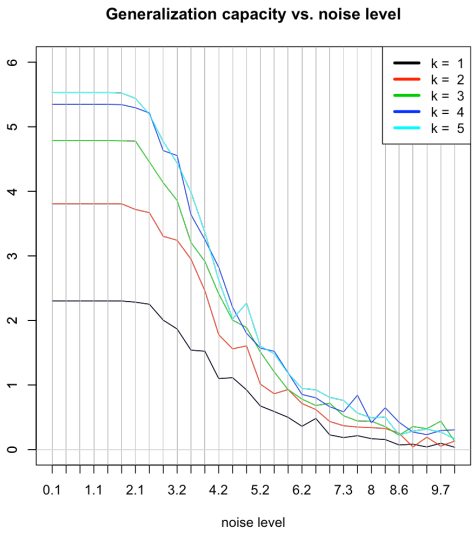

Figure 3.3: Generalization capacity for different k

Again, since values of k closer to $\lceil \frac{d}{2} \rceil$ correspond to larger hypothesis spaces, the generalization capacity of the cost function $\hat{R}(\mu, \xi)$ is larger for these k.

3.3 Sampling algorithm

As was mentioned in Section 1.2 we are ultimately interested in the sparse case where k is kept constant and d grows toward infinity. As d grows, the size of the hypothesis space $\mathcal{C}^k$ grows exponentially fast which means computing the partition functions $Z_\beta(\xi)$ and $\Delta Z_\beta(\xi_1, \xi_2)$, which are sums over $\mathcal{C}^k$, quickly becomes unfeasible. In this section we use the sampling algorithm suggested in **?** to estimate $Z_\beta(\xi)$ and $\Delta Z_\beta(\xi_1, \xi_2)$ without summing over the entire hypothesis space.

Recall that:

$$\mathbb{E}_{(\xi_1,\xi_2)}[I_\beta] = \log|\mathcal{C}^k| + \mathbb{E}_{(\xi_1,\xi_2)}[\log \Delta Z_\beta(\xi_1,\xi_2)] \tag{3.3.0.1}$$
$$- \mathbb{E}_{\xi_1}[\log Z_\beta(\xi_1)] - \mathbb{E}_{\xi_2}[\log Z_\beta(\xi_2)] \tag{3.3.0.2}$$
$$= \log|\mathcal{C}^k| + \mathbb{E}_{(\xi_1,\xi_2)}\left[\log \sum_{\mu\in\mathcal{C}^k} e^{-\beta(\hat{R}(\mu,\xi_1)+\hat{R}(\mu,\xi_2))}\right] \tag{3.3.0.3}$$
$$- \mathbb{E}_{\xi_1}\left[\log \sum_{\mu\in\mathcal{C}^k} e^{-\beta\hat{R}(\mu,\xi_1)}\right] - \mathbb{E}_{\xi_2}\left[\log \sum_{\mu\in\mathcal{C}^k} e^{-\beta\hat{R}(\mu,\xi_2)}\right] \tag{3.3.0.4}$$

If we let $p(\mu) = \frac{1}{|\mathcal{C}^k|} = \frac{1}{\binom{d}{k}}$ then,

$$\mathbb{E}_{(\xi_1,\xi_2)}[I_\beta] = \log|\mathcal{C}^k| + \mathbb{E}_{(\xi_1,\xi_2)}\left[\log|\mathcal{C}^k| \sum_{\mu\in\mathcal{C}^k} \frac{1}{|\mathcal{C}^k|} e^{-\beta(\hat{R}(\mu,\xi_1)+\hat{R}(\mu,\xi_2))}\right] \tag{3.3.0.5}$$
$$- \mathbb{E}_{\xi_1}\left[\log|\mathcal{C}^k| \sum_{\mu\in\mathcal{C}^k} \frac{1}{|\mathcal{C}^k|} e^{-\beta\hat{R}(\mu,\xi_1)}\right] \tag{3.3.0.6}$$
$$- \mathbb{E}_{\xi_2}\left[\log|\mathcal{C}^k| \sum_{\mu\in\mathcal{C}^k} \frac{1}{|\mathcal{C}^k|} e^{-\beta\hat{R}(\mu,\xi_2)}\right] \tag{3.3.0.7}$$
$$= \log|\mathcal{C}^k| + \log|\mathcal{C}^k| + \mathbb{E}_{(\xi_1,\xi_2)}\left[\log \sum_{\mu\in\mathcal{C}^k} \mathbb{E}_\mu\left[e^{-\beta(\hat{R}(\mu,\xi_1)+\hat{R}(\mu,\xi_2))}\right]\right] \tag{3.3.0.8}$$
$$- \log|\mathcal{C}^k| - \mathbb{E}_{\xi_1}\left[\log \sum_{\mu\in\mathcal{C}^k} \mathbb{E}_\mu\left[e^{-\beta\hat{R}(\mu,\xi_1)}\right]\right] \tag{3.3.0.9}$$
$$- \log|\mathcal{C}^k| - \mathbb{E}_{\xi_2}\left[\log \sum_{\mu\in\mathcal{C}^k} \mathbb{E}_\mu\left[e^{-\beta\hat{R}(\mu,\xi_2)}\right]\right] \tag{3.3.0.10}$$
$$= \mathbb{E}_{(\xi_1,\xi_2)}\left[\log \sum_{\mu\in\mathcal{C}^k} \mathbb{E}_\mu\left[e^{-\beta(\hat{R}(\mu,\xi_1)+\hat{R}(\mu,\xi_2))}\right]\right] \tag{3.3.0.11}$$
$$- \mathbb{E}_{\xi_1}\left[\log \sum_{\mu\in\mathcal{C}^k} \mathbb{E}_\mu\left[e^{-\beta\hat{R}(\mu,\xi_1)}\right]\right] \tag{3.3.0.12}$$
$$- \mathbb{E}_{\xi_2}\left[\log \sum_{\mu\in\mathcal{C}^k} \mathbb{E}_\mu\left[e^{-\beta\hat{R}(\mu,\xi_2)}\right]\right] \tag{3.3.0.13}$$

The last expression suggests the following sampling algorithm to estimate the generalization capacity I:

1. Choose a grid of relevant β values: $\underline{\beta} = (\beta_1, ..., \beta_l)$
2. For $i = 1$ to m
   a. Simulate $\xi_1^i, \xi_2^i \sim N(0, I_d)$
   b. Uniformly sample r hypotheses $\mu_j \in \mathcal{C}^k$
   c. For $k = 1$ to l
      - Calculate *quasi* information content:

$$\tilde{I}^i_{\beta_k} = \log \frac{1}{r} \sum_{j=1}^{r} e^{-\beta_k (\hat{R}(\mu_j, \xi_1^i) + \hat{R}(\mu_j, \xi_2^i))} \tag{3.3.0.14}$$

$$- \log \frac{1}{r} \sum_{j=1}^{r} e^{-\beta_k \hat{R}(\mu_j, \xi_1^i)} - \log \frac{1}{r} \sum_{j=1}^{r} e^{-\beta_k \hat{R}(\mu_j, \xi_2^i)} \tag{3.3.0.15}$$

3. For $k = 1$ to l
   - Estimate mean information content:

$$\bar{I}_{\beta_k} = \frac{1}{m} \sum_{i=1}^{m} \tilde{I}^i_{\beta_k} \tag{3.3.0.16}$$

   - Estimate generalization capacity:

$$\hat{I} = \max_{k \in \{1,...,l\}} \bar{I}_{\beta_k} \tag{3.3.0.17}$$

3.4 Simulation results: sampling algorithm

We used the following parameters for the simulation experiments:

- $d = 10$,
- $k = 4$,
- $m = 100$,
- $r = 100, 1000$,
- 20 different β values from 0.01 to 10, and
- 20 different noise levels σ from 0.1 to 15.

We first compare the generalization capacity estimation using the exhaustive and sampling algorithms for two different choices of r.

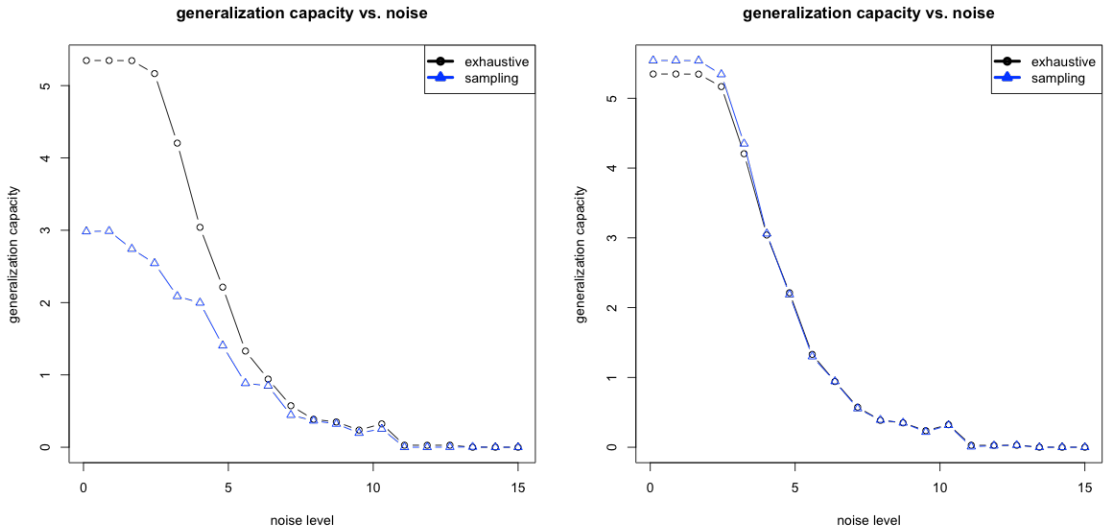

(a) $r = 100$ (b) $r = 1000$

Figure 3.4: Generalization capacity estimation for different r

It is clear that using the sampling algorithm with $r = 100$ we underestimate the generalization capacity since for very low noise levels we know the generalization capacity is equal to $\log |\mathcal{C}^k| = \log |\binom{d}{k}| = \log |\binom{10}{4}| \approx 5.35$. For $r = 1000$ the estimation is much better, however since $|\mathcal{C}^k| = 210$ it is cheaper to use the exhaustive algorithm. Next we check the generalization capacity estimate for both algorithms as d increases. The parameters used for the simulation experiments were the following:

- $d = 5, 6, ..., 20$,
- $k = 4$,
- $m = 100$,
- $r = 20, 100, 500, 1000$,
- 20 different β values from 0.01 to 10, and
- 2 different noise levels $\sigma = 2, 4$.

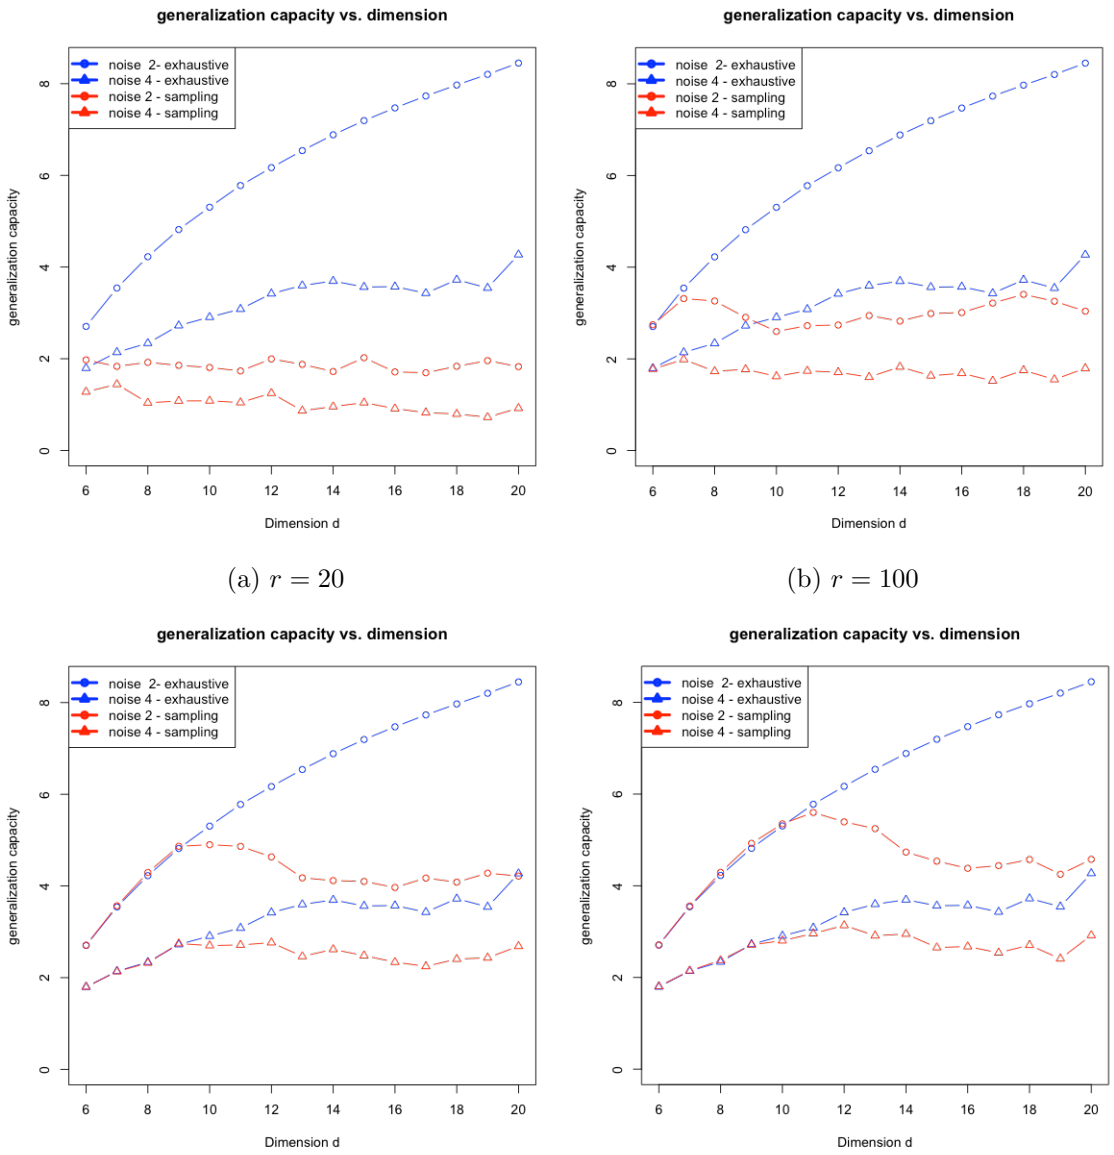

(a) $r = 20$ (b) $r = 100$

(c) $r = 500$ (d) $r = 1000$

Figure 3.5: Generalization capacity estimation for different d and r

The figure above confirms that the sampling algorithm proposed turns out to be more expensive than the exhaustive algorithm. For example, to estimate the generalization capacity accurately when $d = 11$ we need $r = 1000$ which is larger than the size of the hypothesis space $\binom{11}{4} = 330$. The sampling algorithm represents a way to estimate genearalization capacity by summing over a sample of the hypothesis class. The sample size required must be, at most, linear in d if it is to be useful in estimating generalization capacity for sparse conditions when k is fixed and d is large.

3.5 Importance sampling algorithm

The aim of the sampling algorithm is to estimate $\mathbb{E}_\mu[e^{-\beta\hat{R}(\mu,\xi)}]$ with a sample from $p(\mu) = \frac{1}{|\mathcal{C}^k|}$ instead of exhaustively calculating it as:

$$\mathbb{E}_\mu[e^{-\beta\hat{R}(\mu,\xi)}] = \frac{1}{|\mathcal{C}^k|}\sum_{\mu\in\mathcal{C}^k} e^{-\beta\hat{R}(\mu,\xi)} \tag{3.5.0.1}$$

This means sampling $\mu_j \in \mathcal{C}^k$, calculating the Boltzmann weights $w_\beta(\mu_j,\xi)$ and averaging them. In general, different hypothesis $\mu_j \in \mathcal{C}^k$ contribute differently to the average. The lower the cost of a hypothesis the larger the weight contributed. Those hypotheses which contribute most of the weight are relatively small in number. This means that with a small number of samples the proportion of *important* and *unimportant* hypotheses will not accurately reflect the population proportions and the result will be a biased estimate of $Z_\beta(\xi)$ and $\Delta Z_\beta(\xi_1,\xi_2)$. If we take a large sample, of size bigger than the size of $\mathcal{C}^k$, the simulation results seem to show that the estimates converge to those of the exhaustive algorithm, but this defeats the purpose of sampling hypotheses: to obtain an algorithm that is *not* exponential in d as the size of the hypothesis class is.

We want to design an *importance sampling* algorithm where:

1. We sample according to a proposal distribution $q(\mu)$ which assigns more probability to more *important* hypotheses.
2. When estimating the expectation with the average we assign weights $w(\mu)$ to each sample to correct for the fact we sampled according to *wrong* distribution:

$$\mathbb{E}_{p(\mu)}[e^{-\beta\hat{R}(\mu,\xi)}] = \sum_{\mu\in\mathcal{C}^k} p(\mu)e^{-\beta\hat{R}(\mu,\xi)} = \sum_{i=1}^{\binom{d}{k}} p(\mu_i)e^{-\beta\hat{R}(\mu_i,\xi)} \tag{3.5.0.2}$$

$$= \sum_{i=1}^{\binom{d}{k}} \frac{p(\mu_i)}{q(\mu_i)} q(\mu_i)e^{-\beta\hat{R}(\mu_i,\xi)} = \mathbb{E}_{q(\mu)}[w(\mu)e^{-\beta\hat{R}(\mu,\xi)}] \tag{3.5.0.3}$$

Where $w(\mu) = \frac{p(\mu)}{q(\mu)}$.

Recall that we may write:

$$\hat{R}(\mu,\xi) = \mu^T(1 - 2\frac{\sigma}{\sqrt{n}}\xi - 2\mu^0) \tag{3.5.0.4}$$

$$\hat{R}(\mu,\xi_1) + \hat{R}(\mu,\xi_2) = 2\mu^T(1 - \frac{\sigma}{\sqrt{n}}(\xi_1+\xi_2) - 2\mu^0) \tag{3.5.0.5}$$

Observe that $\mu^T 1 = k$ is constant for the case where the hypothesis space is $\mathcal{C}^k$ and not $\mathcal{C} = \mathbb{B}^d$. Also note that $h := \mu^T\mu^0 \in \{\max(0, 2k-d), ..., k\}$ is the number of *hits* of hypothesis μ: the number of components j such that $\mu_j = \mu_j^0 = 1$. With this in mind we can obtain the following equivalent expression for $\hat{R}(\mu,\xi)$:

$$\hat{R}(\mu,\xi) = -2[h + \frac{\sigma}{\sqrt{n}}\mu^T\xi] \tag{3.5.0.6}$$

$$\hat{R}(\mu,\xi_1) + \hat{R}(\mu,\xi_2) = -2[2h + \frac{\sigma}{\sqrt{n}}\mu^T(\xi_1+\xi_2)] \tag{3.5.0.7}$$

Where $h := h(\mu,\mu^0) = \mu^T\mu^0$. This means we may express the Boltzmann weights as:

$$w_\beta(\mu,\xi) = e^{-\beta\hat{R}(\mu,\xi)} = e^{2\beta[h+\frac{\sigma}{\sqrt{n}}\mu^T\xi]} = e^{2\beta h}e^{2\beta\frac{\sigma}{\sqrt{n}}\mu^T\xi} = f_\beta(h)g_\beta(\mu,\xi) \tag{3.5.0.8}$$

The importance of a hypothesis μ is given by the Boltzmann weights $w_\beta(\mu,\xi)$ and these depend on the functions f_β and g_β. We first analyse $g_\beta(\mu,\xi) := e^{2\beta\frac{\sigma}{\sqrt{n}}\mu^T\xi}$. We have that, for $\beta, k > 0$:

$$\xi \sim N(0,I_d) \Rightarrow \mu^T\xi \sim N(0,k) \Rightarrow \frac{2\beta\sigma}{\sqrt{n}}\mu^T\xi \sim N\left(0,\left(\frac{2\beta\sigma}{\sqrt{n}}\right)^2 k\right) \tag{3.5.0.9}$$

$$\Rightarrow e^{\frac{2\beta\sigma}{\sqrt{n}}\mu^T\xi} \sim \log N\left(0,\left(\frac{2\beta\sigma}{\sqrt{n}}\right)^2 k\right) \tag{3.5.0.10}$$

$$\Rightarrow \mathbb{E}_\xi[e^{\frac{2\beta\sigma}{\sqrt{n}}\mu^T\xi}] = e^{\frac{1}{2}(\frac{2\beta\sigma}{\sqrt{n}})^2 k} > 1 \tag{3.5.0.11}$$

This means that we can expect $g_\beta(\mu,\xi)$ to increase $w_\beta(\mu,\xi)$ for half the samples ξ and decrease it for the other half, although in general, we can expect $g_\beta(\mu,\xi)$ to increase $w_\beta(\mu,\xi)$. So we see that $g_\beta(\mu,\xi)$ may contribute to a hypothesis being more or less important depending on the ξ sampled. If we want to take this effect into account in determining the proposal distribution $q(\mu)$ we would have to make it depend on the given ξ sampled. For simplicity we only take into account $f_\beta(h)$ in determining $q(\mu)$.

We now analyse $f_\beta(h) = e^{2\beta h}$. The more hits h that a hypothesis μ has with respect to μ^0 the lower the costs $\hat{R}(\mu,\xi)$ which means the larger he weight $w_\beta(\mu,\xi)$ contributed. Also note that there are $\binom{k}{h}\binom{d-k}{k-h}$ hypotheses $\mu \in \mathcal{C}^k$ such that $h(\mu,\mu^0) = \mu^T\mu^0 = h$. The following figure illustrates the importance $f_\beta(h)$ of a hypothesis $\mu \in \mathcal{C}^k$ and the number of such hypotheses in $\mathcal{C}^k$ as a function of h for $d = 20$, $k = 4$ and $\beta = 1$.

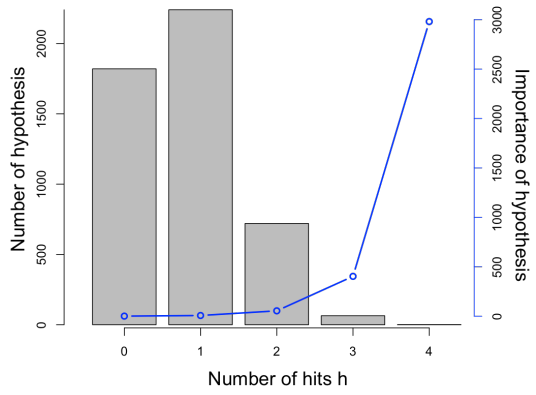

Figure 3.6: Importance and frequency of hypotheses

Hypothesis μ such that $h(\mu, \mu^0)$ is close to k contribute moste of the weight but are relatively few in number. There are $\binom{k}{h}\binom{d-k}{k-h}$ hypotheses $\mu \in \mathcal{C}^k$ such that $h(\mu, \mu^0) = \mu^T\mu^0 = h$ so there are $\sum_{h=\max(1,2k-d)}^{k} \binom{k}{h}\binom{d-k}{k-h}$ hypotheses such that $h > 0$. As d grows this becomes a low proportion of $\binom{d}{k}$, the total number of hypotheses. In other words, as d grows, the most important hypotheses, those that contribute most of the weight, become a smaller proportion of all hypotheses.

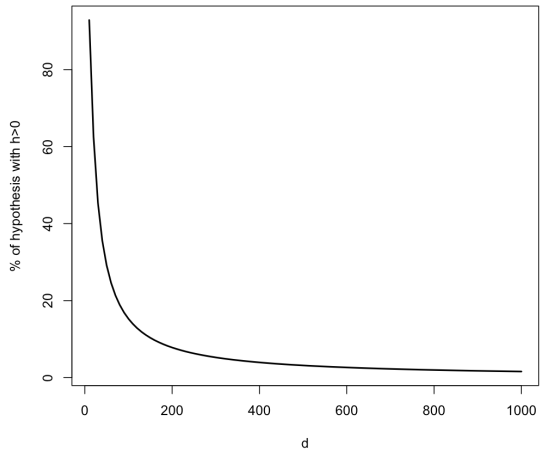

Figure 3.7: % of hypotheses $\mu \in \mathcal{C}^k$ such that $h > 0$

This means that to sample a representative proportion of important hypotheses we need a very large sample. To find a way around this we use a proposal distribution $q(\mu)$ such that the probability of sampling a hypothesis with h hits is the same for all h.

Let,

$$h \sim U\{\max(0, 2k-d), ..., k\} \Rightarrow q(h) = \frac{1}{k - \max(0, 2k-d) + 1} \tag{3.5.0.12}$$

$$\mu|h \sim U\left\{1, ..., \binom{k}{h}\binom{d-k}{k-h}\right\} \Rightarrow q(\mu|h) = \frac{\mathbb{1}_{\{\mu^T\mu^0 = h\}}}{\binom{k}{h}\binom{d-k}{k-h}} \tag{3.5.0.13}$$

This means that:

$$q(\mu) = \sum_{h=\max(0,2k-d)}^{k} q(\mu, h) = \sum_{h=\max(0,2k-d)}^{k} q(\mu|h)q(h) \tag{3.5.0.14}$$

$$= \sum_{h=\max(0,2k-d)}^{k} \frac{\mathbb{1}_{\{\mu^T\mu^0 = h\}}}{(k - \max(0, 2k-d) + 1)\binom{k}{h}\binom{d-k}{k-h}} \tag{3.5.0.15}$$

$$= \frac{1}{(k - \max(0, 2k-d) + 1)\binom{k}{\mu^T\mu^0}\binom{d-k}{k-\mu^T\mu^0}} \tag{3.5.0.16}$$

$$w(\mu) = \frac{p(\mu)}{q(\mu)} = \frac{\binom{k}{\mu^T\mu^0}\binom{d-k}{k-\mu^T\mu^0}}{\binom{d}{k}} \tag{3.5.0.17}$$

The above suggests the following importance sampling algorithm for estimating generalization capacity.

I. Choose a grid of relevant β values: $\underline{\beta} = (\beta_1, ..., \beta_l)$

II. For $i = 1$ to m

1. Simulate $\xi_1^i, \xi_2^i \sim N(0, I_d)$
2. For $j = 1$ to r

   a. Sample from $h_j \sim q(h) = \frac{1}{k - \max(0, 2k-d) + 1}$

   b. Sample from $\mu_j | h_j \sim q(\mu|h) = \frac{\mathbb{1}_{\{\mu^T\mu^0 = h\}}}{\binom{k}{h}\binom{d-k}{k-h}}$ This can be done by:

      i. Calculate number of *misses* $m_j = k - h_j$

      ii. Identify $A_0 = \{a : \mu_a^0 = 0\}$ and $A_1 = \{a : \mu_a^0 = 1\}$

      iii. Set $\mu_j \leftarrow \mu^0$

      iv. Uniformly sample m_j times $r \in A_0$ and $s \in A_1$ setting $\mu_{jr} = 1$ and $\mu_{js} = 0$ each time.
3. For $k = 1$ to l

   - Calculate *quasi* information content:

$$\tilde{I}_{\beta_k}^i = \log \frac{1}{r} \sum_{j=1}^{r} e^{-\beta_k(\hat{R}(\mu_j, \xi_1^i) + \hat{R}(\mu_j, \xi_2^i))} \tag{3.5.0.18}$$

$$- \log \frac{1}{r} \sum_{j=1}^{r} e^{-\beta_k \hat{R}(\mu_j, \xi_1^i)} - \log \frac{1}{r} \sum_{j=1}^{r} e^{-\beta_k \hat{R}(\mu_j, \xi_2^i)} \tag{3.5.0.19}$$

III. For $k = 1$ to l

- Estimate mean information content:

$$\bar{I}_{\beta_k} = \sum_{i=1}^{m} w_i \tilde{I}_{\beta_k}^i \tag{3.5.0.20}$$

  Where $w_i = w(\mu_i) = \frac{\binom{k}{\mu_i^T\mu^0}\binom{d-k}{k-\mu_i^T\mu^0}}{\binom{d}{k}}$

- Estimate generalization capacity:

$$\hat{I} = \max_{k \in \{1,...,l\}} \bar{I}_{\beta_k} \tag{3.5.0.21}$$

3.6 Simulation results: importance sampling algorithm

We used the following parameters for the simulation experiments:

- $d = 10$,
- $k = 4$,
- $m = 100$,
- $r = 100, 1000$,
- 20 different β values from 0.01 to 10, and
- 20 different noise levels σ from 0.1 to 15.

We first compare the generalization capacity estimation using the exhaustive, sampling and importance sampling algorithms for two different choices of r.

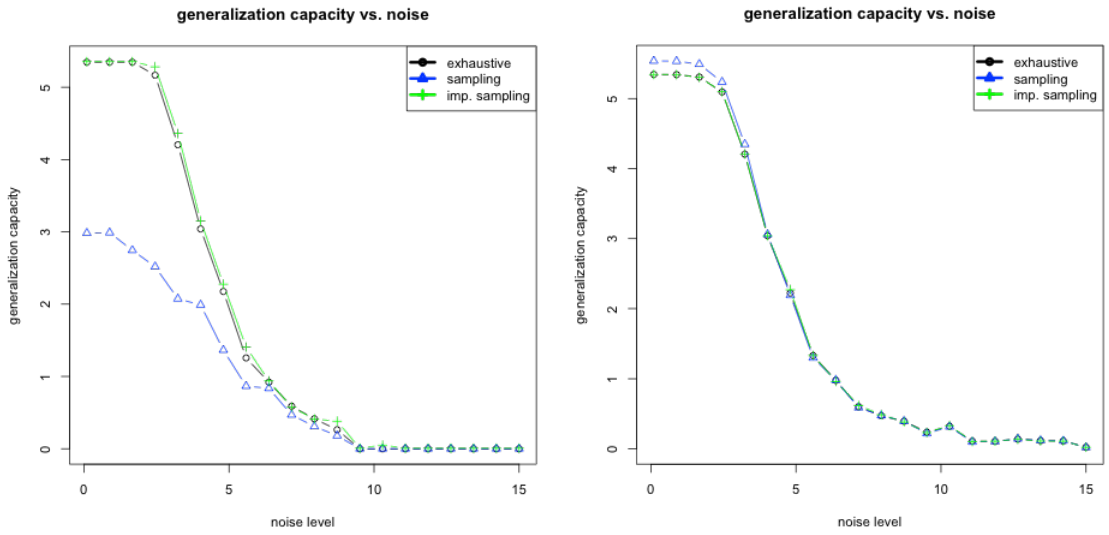

(a) $r = 100$ (b) $r = 1000$

Figure 3.8: Generalization capacity estimation for different r

The figures suggest that the generalization capacity estimated with the importance sampling algorithm converges to the exhaustive algorithm results for a sample size much smaller than is required for the sampling algorithm. Next we check the generalization capacity estimate for all three algorithms as d increases. The parameters used for the simulation experiments were the following:

- $d = 5, 6, ..., 20$,
- $k = 4$,
- $m = 100$,
- $r = 20, 100, 500, 1000$,
- 20 different β values from 0.01 to 10, and

- 2 different noise levels $\sigma = 2, 4$.

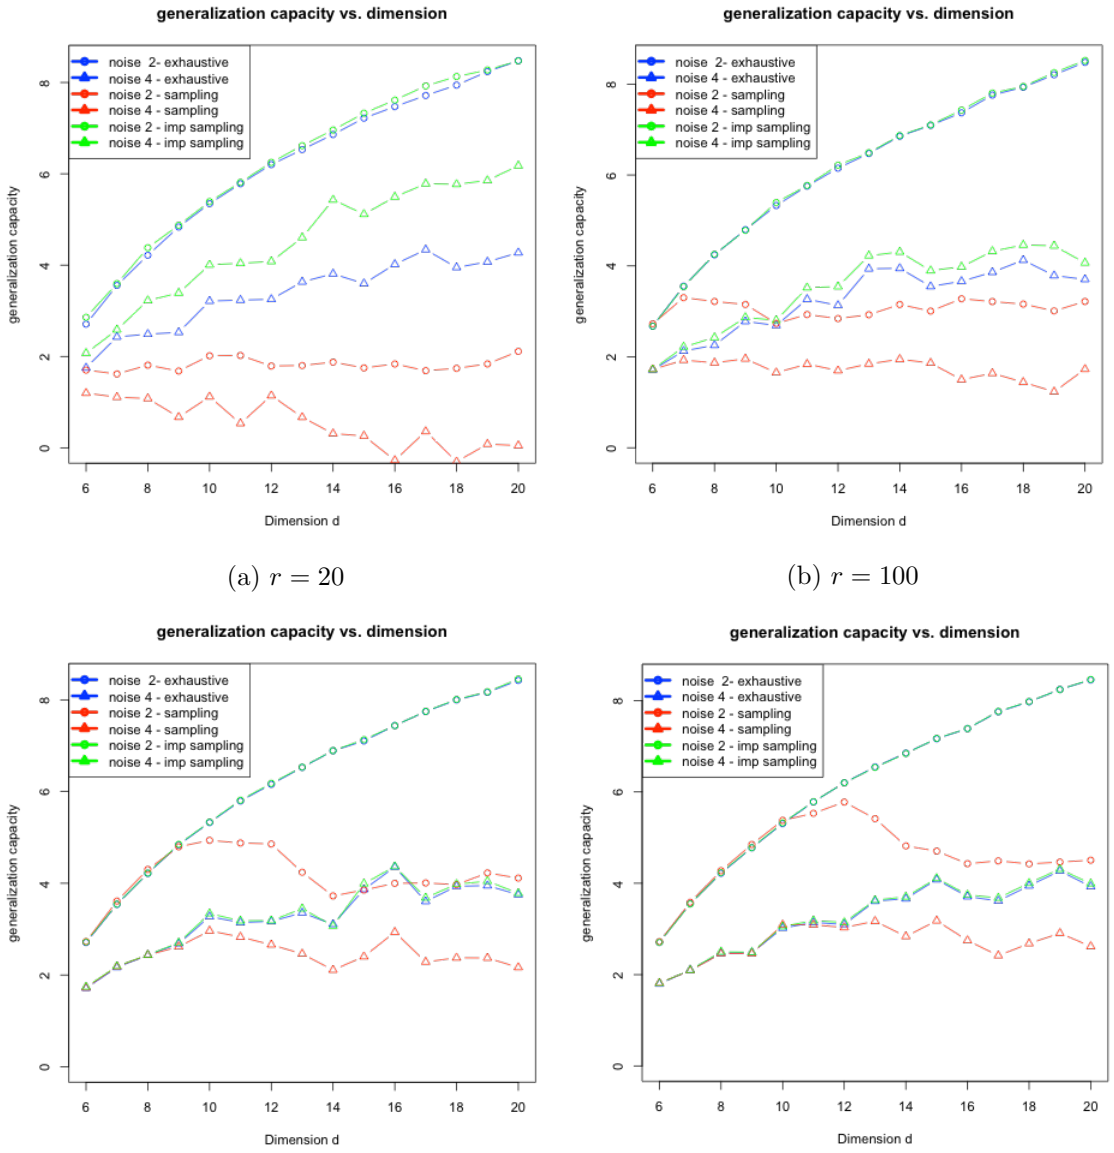

(a) $r = 20$ (b) $r = 100$

(c) $r = 500$ (d) $r = 1000$

Figure 3.9: Generalization capacity estimation for different d and r

The above figures seem to confirm that the importance sampling algorithm need a much smaller sample size to converge than does the sampling algorithm suggesting it will be useful in estimating generalization capacity for large values of d when we can no longer *exhaustively* evaluate the partition functions. In the following simulation experiments we estimate the generalization capacity of the cost function $\hat{R}(\mu, \xi)$ in sparse settings where k is constant and d is large. We used the following parameters for these simulation experiments:

- $d = 10, ..., 10000$,
- $k = 4$,

- $m = 100$,
- $r = 100$,
- 100 different β values from 0.01 to 30, and
- 4 different noise levels $\sigma = 3, 4, 5, 6$.

Recall we ended Section 1.6 by mentioning that we can use generalization capacity to choose between alternative cost functions. The following simulation experiments, carried out with the above parameters, were performed for two different cost functions, the squared-loss based risk function and an absolute-loss based risk function. Recall expression 2.1.0.8 for the squared-loss based risk function:

$$\hat{R}(\mu, X) = ||\mu - \bar{X}||_2^2 \tag{3.6.0.1}$$

If we replace the L2 norm with the L1 norm we get te absolute-loss based risk function:

$$\hat{R}(\mu, X) = ||\mu - \bar{X}||_1^2 \tag{3.6.0.2}$$

The following figure shows the estimated generalization capacity of both cost functions for different number of components d and noise levels σ. Each scatter point represents the estimation, by simulation, of the generalization capacity for a given d and cost function $\hat{R}(\mu, \xi)$. Blue points represent the generalization capacity of the L1 norm cost function and black points that of the L2 norm cost function. The blue and black lines are smoothing splines applied to the blue and black points respectively.

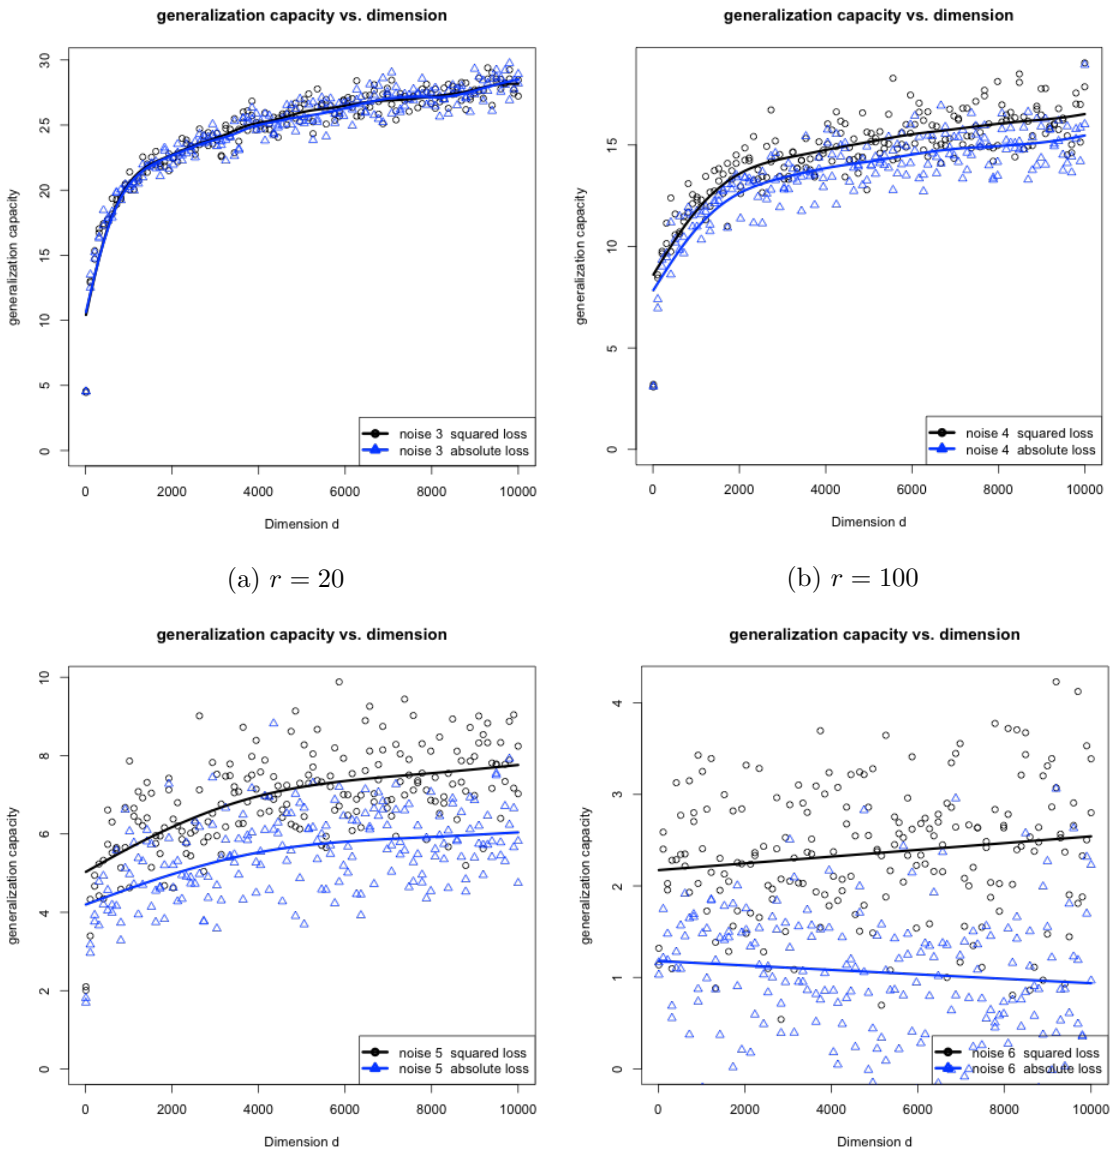

(a) $r = 20$ (b) $r = 100$

(c) $r = 500$ (d) $r = 1000$

Figure 3.10: Generalization capacity estimation for different d and σ

We can see that the generalization capacity for both cost functions, as a function of the number of components d, displays concave behavior. Initially, as the hypothesis set size grows we see rapid increase in the generalization capacity however after a certain threshold, the large number of hypotheses makes detection of the *relevant* components difficult so that additional gains in generalization capacity are marginal. The value of this threshold depends on the noise level of the data: the higher the noise level the lower the threshold. Additionally, observe that, as we would expect given that the data has a gaussian distribution, the L2 norm cost function has a greater generalization capacity than that of the L1 norm cost function.

Chapter 4

Towards sparse feature selection for correlated features

As with the model described in Section 1.2 we deal with the statistical model:

$$X_i = \mu^0 + \epsilon_i$$

Except in this case we have that:

- $||\mu^0||_1 = k$,
- $\mu^0 \in \mathbb{B}_k^d = \{\mu \in \mathbb{B}^d : ||\mu||_1 = k\}$,
- $|\mathbb{B}_k^d| = \binom{d}{k}$,
- It is assumed that k is known,
- $\epsilon \sim N(0, \Gamma)$ where $\Gamma_{ii} = \sigma^2$ for all $i \in \{1, ..., d\}$, and
- observations X_i with $i = 1, ..., n$ are i.i.d.

Notice that in contrast to the problem statement in section 1.2, the features are now correlated. This is because each entry represents a variable or feature and the normal situation is that features are correlated. We now change the sparsity condition from a small, fixed k and large d to $k \approx \frac{d}{\log(d)}$. This reflects the assumption that the number of relevant features grows as the number of total features grows, albeit at a slower rate.

We are ultimately interested in estimating, by simulation, the generalization capacity of the quadratic loss based, empirical risk function. We could try to use the importance sampling algorithm of Section 3.5 although we haven't tested for the case where $k \approx \frac{d}{\log d}$ and $\Sigma \neq I_d$. In this Chapter we explore an alternate way to approximate the expected value of the log partition function $\mathbb{E}_X[\log Z_\beta(X)]$, which, as we have seen, is a component of the generalization capacity. The following derivations and approximations are based on **?**.

In this case we have that

$$X \sim N\Big(\mu^0, \Gamma\Big) \Rightarrow \bar{X} \sim N\Big(\mu^0, \frac{\Gamma}{n}\Big) \tag{4.0.0.1}$$

So we let $\xi = \frac{\bar{X}-\mu^0}{\sigma/\sqrt{n}} \sim N(0,\Sigma)$, where $\Sigma = \frac{\Gamma}{\sigma^2}$ so that $\bar{X} = \mu^0 + \frac{\sigma}{\sqrt{n}}\xi$. Recall from 2.5.0.26 that we may write the square based empirical risk function as:

$$\hat{R}(\mu,\xi) = \mu^T(1 - 2\frac{\sigma}{\sqrt{n}}\xi - 2\mu^0) \tag{4.0.0.2}$$

$$\tag{4.0.0.3}$$

Let $\hat{r}(\xi) = 1 - 2\frac{\sigma}{\sqrt{n}}\xi - 2\mu^0$ and $\mathcal{C}^k_h = \{\mu \in \mathcal{C}^k : \mu^T\mu^0 = h\}$. Then:

$$\hat{R}(\mu,\xi) = \mu^T\hat{r}(\xi) \tag{4.0.0.4}$$

$$Z_\beta(\xi) = \sum_{\mu\in\mathcal{C}^k} w_\beta(\mu,\xi) = \sum_{\mu\in\mathcal{C}^k} e^{-\beta\hat{R}(\mu,\xi)} \tag{4.0.0.5}$$

$$= \sum_{\mu\in\mathcal{C}^k} e^{-\beta\mu^T\hat{r}(\xi)} = \sum_{h=max(0,2k-d)}^{k} \sum_{\mu\in\mathcal{C}^k_h} e^{-\beta\mu^T\hat{r}(\xi)} \tag{4.0.0.6}$$

with which

$$\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big] = \mathbb{E}_\xi\Big[\log \sum_{h=max(0,2k-d)}^{k} \sum_{\mu\in\mathcal{C}^k_h} e^{-\beta\mu^T\hat{r}(\xi)}\Big] \tag{4.0.0.7}$$

$$= \mathbb{E}_\xi\Big[\log\Big\{\sum_{\mu\in\mathcal{C}^k_k} e^{-\beta\mu^T\hat{r}(\xi)} + \sum_{h=max(0,2k-d)}^{k-1} \sum_{\mu\in\mathcal{C}^k_h} e^{-\beta\mu^T\hat{r}(\xi)}\Big\}\Big] \tag{4.0.0.8}$$

but since $\mathcal{C}^k_k = \{\mu^0\}$ we have that

$$\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big] = \mathbb{E}_\xi\Big[\log\Big\{e^{-\beta\mu^{0T}\hat{r}(\xi)} + \sum_{h=max(0,2k-d)}^{k-1}\binom{k}{h}\binom{d-k}{k-h}\sum_{\mu\in\mathcal{C}^k_h}\frac{e^{-\beta\mu^T\hat{r}(\xi)}}{\binom{k}{h}\binom{d-k}{k-h}}\Big\}\Big] \tag{4.0.0.9}$$

$$= \mathbb{E}_\xi\Big[\log\Big\{e^{-\beta\mu^{0T}\hat{r}(\xi)}\Big(1 + e^{\beta\mu^{0T}\hat{r}(\xi)}\sum_{h=max(0,2k-d)}^{k-1}\binom{k}{h}\binom{d-k}{k-h}\sum_{\mu\in\mathcal{C}^k_h}\frac{e^{-\beta\mu^T\hat{r}(\xi)}}{\binom{k}{h}\binom{d-k}{k-h}}\Big)\Big\}\Big] \tag{4.0.0.10}$$

$$= \mathbb{E}_\xi\Big[-\beta\mu^{0T}\hat{r}(\xi)\Big] + \mathbb{E}_\xi\Big[\log\Big\{1 + \sum_{h=max(0,2k-d)}^{k-1}\binom{k}{h}\binom{d-k}{k-h}\sum_{\mu\in\mathcal{C}^k_h}\frac{e^{-\beta(\mu-\mu^0)^T\hat{r}(\xi)}}{\binom{k}{h}\binom{d-k}{k-h}}\Big\}\Big] \tag{4.0.0.11}$$

Let

$$\bar{B}_h := \sum_{\mu \in \mathcal{C}_h^k} \frac{e^{-\beta(\mu-\mu^0)^T \hat{r}(\xi)}}{\binom{k}{h}\binom{d-k}{k-h}} \tag{4.0.0.12}$$

Notice that $\bar{B}_k = 1$ so that

$$\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big] = -\beta\mu^{0T}\mathbb{E}_\xi\Big[\hat{r}(\xi)\Big] + \mathbb{E}_\xi\Big[\log \sum_{h=max(0,2k-d)}^{k} \binom{k}{h}\binom{d-k}{k-h}\bar{B}_h\Big] \tag{4.0.0.13}$$

If we take a hypothesis μ at random the number of hits $h := h(\mu)$ is a hypergeometric random variable:

$$H \sim Hypergeometric(S = k, N = d, n = k) \tag{4.0.0.14}$$

Where

- H measures the number of succesesses from n draws without replacement,
- S is the number of success states and
- n is the number of draws.

In this case we pick a μ at random which represents picking k components to be equal to one (without replacement), from a population of d where there are exactly k components equal to one. This random variable has the following probability distribution function:

$$\mathbb{P}(H = h) = \frac{\binom{k}{h}\binom{d-k}{k-h}}{\binom{d}{k}} \tag{4.0.0.15}$$

Where $h \in \{\max(0, n + S - N), ..., \min(n, S)\} = \{\max(0, 2k - d), ..., k\}$ which means that

$$\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big] = -\beta\mu^{0T}\mathbb{E}_\xi\Big[\hat{r}(\xi)\Big] + \mathbb{E}_\xi\Big[\log \binom{d}{k} \sum_{h=\max(0,2k-d)}^{k} \frac{\binom{k}{h}\binom{d-k}{k-h}}{\binom{d}{k}}\bar{B}_h\Big] \tag{4.0.0.16}$$

$$= -\beta\mu^{0T}\mathbb{E}_\xi\Big[\hat{r}(\xi)\Big] + \mathbb{E}_\xi\Big[\log \binom{d}{k} \sum_{h=\max(0,2k-d)}^{k} \mathbb{P}(H = h)\bar{B}_h\Big] \tag{4.0.0.17}$$

$$= -\beta\mu^{0T}\mathbb{E}_\xi\Big[\hat{r}(\xi)\Big] + \mathbb{E}_\xi\Big[\log \binom{d}{k}\mathbb{E}_H\Big[\bar{B}_H\Big]\Big] \tag{4.0.0.18}$$

Now $|\mathcal{C}_h^k|$ is exponential in d and $\bar{B}_h$ is a sum over a very large hypothesis space so we will need a way to approximate it. We have that:

$$\mu^T\hat{r}(\xi) = \mu^T(1 - 2\mu^0 - 2\frac{\sigma}{\sqrt{n}}) = \mu^T 1 - 2\mu^T\mu^0 - 2\frac{\sigma}{\sqrt{n}}\mu^T\xi \tag{4.0.0.19}$$

$$= k - 2h - \eta(\mu) \tag{4.0.0.20}$$

where $\eta(\mu) := 2\frac{\sigma}{\sqrt{n}}\mu^T\xi$. First notice that if $\mu = \mu^0$ then $h = k$, $\mu^T\hat{r}(\xi) = -k - \eta(\mu)$ and since $\xi \sim N(0, \Sigma)$ we have that $\mu^T\mathbb{E}_\xi\Big[\hat{r}(\xi)\Big] = -k$ so that:

$$\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big] = \beta k + \mathbb{E}_\xi\Big[\log\binom{d}{k}\mathbb{E}_H\Big[\bar{B}_H\Big]\Big] \tag{4.0.0.21}$$

$$= \beta k + \mathbb{E}_\xi\Big[\log\binom{d}{k}\mathbb{E}_H\Big[\sum_{\mu\in\mathcal{C}_H^k}\frac{e^{-\beta(\mu-\mu^0)^T\hat{r}(\xi)}}{\binom{k}{H}\binom{d-k}{k-H}}\Big]\Big] \tag{4.0.0.22}$$

$$= \beta k + \mathbb{E}_\xi\Big[\log\binom{d}{k}\mathbb{E}_H\Big[\sum_{\mu\in\mathcal{C}_H^k}\frac{e^{-\beta\{(k-2H-\eta(\mu))-(-k-\eta(\mu^0))\}}}{\binom{k}{H}\binom{d-k}{k-H}}\Big]\Big] \tag{4.0.0.23}$$

$$= \beta k + \mathbb{E}_\xi\Big[\log\binom{d}{k}e^{-2\beta k}e^{-\beta\eta(\mu^0)}\mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}}\sum_{\mu\in\mathcal{C}_H^k}e^{\beta\eta(\mu)}\Big]\Big] \tag{4.0.0.24}$$

$$= \beta k + \mathbb{E}_\xi\Big[\log\binom{d}{k} - 2\beta k - \beta\eta(\mu^0) + \log\mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}}Y_H(\xi)\Big]\Big] \tag{4.0.0.25}$$

$$= -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[-\beta\eta(\mu^0)\Big] + \mathbb{E}_\xi\Big[\log\mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}}Y_H(\xi)\Big]\Big] \tag{4.0.0.26}$$

Where $Y_h := \sum_{\mu\in\mathcal{C}_h^k} e^{\beta\eta(\mu)}$. Since

$$\mathbb{E}_\xi\Big[-\beta\eta(\mu^0)\Big] = \mathbb{E}_\xi\Big[-2\beta\frac{\sigma}{\sqrt{n}}\mu^{0T}\xi\Big] = -2\beta\frac{\sigma}{\sqrt{n}}\mu^{0T}\mathbb{E}_\xi[\xi] = 0 \tag{4.0.0.27}$$

so that

$$\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big] = -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[\log\mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}}Y_H(\xi)\Big]\Big] \tag{4.0.0.28}$$

The problem with evaluating Y_h is that it is a sum over $\mathcal{C}_h^k$ which is very large if d is large and $k \approx \frac{d}{\log d}$. If we can come up with a good approximation $\widetilde{Y}_h$ of Y_h then the above formula suggests the following algorithm for estimating $\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big]$, for a given β, d and k:

1. For $i = 1$ to m

a. Simulate $\xi^i \sim N(0, \Sigma)$

b. For $j = 1$ to p

- Simulate $H_j \sim Hypergeometric(k, d, k)$
- Calculate a_j

$$a_j := \frac{e^{2\beta H_j}}{\binom{k}{H_j}\binom{d-k}{k-H_j}} \widetilde{Y}_{H_j}(\xi^i) \tag{4.0.0.29}$$

c. Calculate $\bar{a}^i = \frac{1}{p}\sum_{j=1}^{p} a_j$

2. Estimate $\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big]$

$$\hat{\mathbb{E}}_\xi\Big[\log Z_\beta(\xi)\Big] = -\beta k + \log\binom{d}{k} + \frac{1}{m}\sum_{i=1}^{m} \log \bar{a}^i \tag{4.0.0.30}$$

In order to explore possible ways to approximate Y_h we take a look at its distribution.

$$Y_h = \sum_{\mu \in \mathcal{C}_h^k} e^{\beta\eta(\mu)} = \sum_{\mu \in \mathcal{C}_h^k} e^{2\beta\frac{\sigma}{\sqrt{n}}\mu^T\xi} \tag{4.0.0.31}$$

Since

$$\xi \sim N(0, \Sigma) \Rightarrow \mu^T\xi \sim N(0, \mu^T\Sigma\mu) \tag{4.0.0.32}$$

$$\Rightarrow \eta(\mu) = \frac{2\sigma}{\sqrt{n}}\mu^T\xi \sim N\Big(0, \frac{4\sigma^2}{n}\mu^T\Sigma\mu\Big) \tag{4.0.0.33}$$

$$\Rightarrow e^{\beta\eta(\mu)} \sim \log N\Big(0, \frac{4\beta^2\sigma^2}{n}\mu^T\Sigma\mu\Big) \tag{4.0.0.34}$$

and

$$Cov(\mu_1^T\xi, \mu_2^T\xi) = \mu_1^T\Sigma\mu_2 \Rightarrow Cov(\beta\eta(\mu_1), \beta\eta(\mu_2)) = \frac{4\beta^2\sigma^2}{n}\mu_1^T\Sigma\mu_2 \tag{4.0.0.35}$$

we have that even in the case that $\Sigma = I_d$, $Y_h(\xi)$ is a sum of $|\mathcal{C}_h^k|$ correlated log-normals with correlation matrix Λ such that

$$\Lambda_{ij} = \frac{4\beta^2\sigma^2}{n}\mu_i^T\Sigma\mu_j \tag{4.0.0.36}$$

If $\Sigma = I_d$ we have that

$$\beta\eta(\mu) \sim N\left(0, \frac{4\beta^2\sigma^2 k}{n}\right) \tag{4.0.0.37}$$

$$\Rightarrow e^{\beta\eta(\mu)} \sim \log N\left(0, \frac{4\beta^2\sigma^2 k}{n}\right) \tag{4.0.0.38}$$

and

$$Cov(\beta\eta(\mu_i), \beta\eta(\mu_j)) = \Lambda_{ij} = \frac{4\beta^2\sigma^2 c_{ij}}{n} \tag{4.0.0.39}$$

$$Cor(\beta\eta(\mu_i), \beta\eta(\mu_j)) = \frac{\Lambda_{ij}}{\sqrt{\Lambda_{ii}}\sqrt{\Lambda_{jj}}} = \frac{c_{ij}}{k} \tag{4.0.0.40}$$

Where $c_{ij} := \mu_i^T \mu_j$ is the number of overlapping ones in μ_i and μ_j. If we can use the approximation proposed in **?** where we assume $\eta(\mu) \approx \eta_h$ we have that:

$$\mathbb{E}_\xi\Big[\log Z_\beta(\xi)\Big] = -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[\log \mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}} Y_H(\xi)\Big]\Big] \tag{4.0.0.41}$$

$$= -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[\log \mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}} \sum_{\mu\in\mathcal{C}_H^k} e^{\beta\eta(\mu)}\Big]\Big] \tag{4.0.0.42}$$

$$\approx -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[\log \mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}} \sum_{\mu\in\mathcal{C}_H^k} e^{\beta\eta_H}\Big]\Big] \tag{4.0.0.43}$$

$$= -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[\log \mathbb{E}_H\Big[\frac{e^{2\beta H}}{\binom{k}{H}\binom{d-k}{k-H}} |\mathcal{C}_H^k| e^{\beta\eta_H}\Big]\Big] \tag{4.0.0.44}$$

$$= -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[\log \mathbb{E}_H\Big[e^{2\beta H} e^{\beta\eta_H}\Big]\Big] \tag{4.0.0.45}$$

$$= -\beta k + \log\binom{d}{k} + \mathbb{E}_\xi\Big[\log \mathbb{E}_H\Big[e^{\beta(2H+\eta_H)}\Big]\Big] \tag{4.0.0.46}$$

However it is not clear that this approximation will be helpful since it corresponds to assuming that if $\mu_i, \mu_j \in \mathcal{C}_{h1}^k$ and $\mu_p \in \mathcal{C}_{h2}^k$ then $Cor(\beta\eta(\mu_i), \beta\eta(\mu_j)) = 1$ and $Cor(\beta\eta(\mu_i), \beta\eta(\mu_p)) = 0$ for all $i, j \in \{1, ..., \binom{k}{h1}\binom{d-k}{k-h1}\}$ and $p \in \{1, ..., \binom{k}{h2}\binom{d-k}{k-h2}\}$.

Chapter 5

Summary

The work presented in this report falls into the following three categories:

i.) Exposition of the theory related to approximation set coding and generalization capacity: Sections 1.3-1.6,

ii.) Implementation of the generalization capacity concept to the sparse mean localization: Section 1.2 and Chapters 2-3, and

iii.) Exploration of how generalization capacity can be implemented for a more realistic version of the sparse feature selection problem: Chapter 4.

In the first case we introduce pattern analysis which is a broad framework with which to deal with learning problems. We describe the approximation set approach to learning where we look for a set of hypotheses with *similar* performance instead of looking for just one. With the definitions from pattern analysis and approximation sets we then introduce information theoretic concepts, originally developed within the field of statistical physics, which we need to define GC. To motivate the meaning of GC we study Shannon's noisy channel coding theorem and related theory drawing an extensive analogy between channel capacity and generalization capacity. We describe in detail the differences and similarities between the communication scenario from which Shannon's concept of channel capacity arises and the learning scenario from which **?**'s concept of generalization capacity can be derived.

The second type of contribution presented here involves the implementation of generalization capacity to the problem of sparse mean localization. In **?** an expression for GC for the sparse mean localization problem (with respect to the squared loss based empirical risk function) is derived that suggests a simulating algorithm for its estimation. The simulating algorithm is briefly described and results are presented. Based on this work we designed and detailed successive simulating algorithms for GC estimation for non-sparse and sparse mean localization problems. We first designed an *exhaustive* sampling algorithm for non-sparse mean localization in low dimensional cases where the dimension d of the mean vector μ^0 is in the order of 10. We implemented various versions of *common random numbers* to reduce variance and chose the best one. We then applied it to the sparse mean localization problem and tried to scale it up to deal with high dimensional cases, but found that since the size of the hypothesis class is exponential in the dimension d the exhaustive algorithm is unfeasible since it involves summing over the entire hypothesis space. To deal with this we designed a uniform sampling algorithm, also suggested in **?**,

but found that for this algorithm to converge the number of sampled hypotheses necessary was actually larger than the hypothesis class size. We solved this difficulty by designing an importance sampling alogorithm which seems to be adequate for high dimensions in the order of 10,000 (asymptotic confidence intrevals need to be derived to verify this). In **?** results showing the estimation of GC for high dimensional sparse mean localization are shown however the simulation method is not detailed. The contribution of this work is to detail and justify numerical estimation of GC for the sparse mean localization problem. It is also worth mentioning that the results of the estimatation of GC are qualitatively different to those found in **?**:compare for example figures 1 and 2 in **?** to figure 3.10.

Finally, based on **?**, we explored estimating GC in a more realistic version of the sparse feature selection problem and sketched a general simulation algorithm, describing certain difficulties that need to be resolved.

5.1 Future work

Possible ways to extend the work presented here are:

- Derive asymptotic confidence intervals for the different GC estimators (exhaustive, sampling, importance sampling) so as to verify when these are fullfilled and so determine the number of simulations necessary. This is especially important for the sparse high dimensional cases (d large) so as to establish the reliability of the estimates produced and the feasability of the algorithms.
- Explore the sparse feature selection problem further: does the $\eta(\mu) \approx \eta_h$ approximation work? What other approximations can we use to make the estimation of $Y_h(\xi)$ feasible for large d and $k \approx \frac{d}{\log d}$.
- Implement the importance sampling algorithm to the more general sparse feature selection problem and compare with simulation methods based on approximations of the type discussed in Chapter 4.
- Extend the study to include estimation of sparse histograms. In this case a histogram can be represented as a vector X_i composed of a mean μ^0 and noise ϵ: i.e. $X_i = \mu^0 + \epsilon_i$ as in the sparse mean localization problem. However in this case $\mu^0 \in [0,1]^d$ instead of $\mu^0 \in \{0,1\}^d$ and $|\mu^0|_1 = 1$ instead of $|\mu^0|_1 = k$. The number of non-zero entries is still assumed to be k to enforce a sparsity condition.

Appendices

Appendix A

R Code

We include the code in R that was used to implement the importance sampling algorithm. Other algorithms described to estimate generalization capacity are very similar so we omit them.

A.1 Functions

```
##############################################################################
# Aproximation Set Coding
# Sparse Mean localization
# Importance Sampling Algorithm
##############################################################################

##############################################################################
# Functions
##############################################################################

#Generate any of the 1,...,choose(d,k) hypothesis
get.hyp.i <- function(i,d,k){

  one.pos <- as.numeric()
  m <- i-1

  for(l in k:1){
    candidates <- d:(l-1)
    try.ck <- choose(candidates,l)
    indx.ck <- which(try.ck <= m)[1]
    ck <- candidates[indx.ck]
    one.pos <- c(one.pos,ck)
    m <- m - choose(ck,l)

  }
  res <- rep(0,d)
  res[one.pos+1] <-1
  return(res)
}

#Generate unsrestricted hypothesis space
get.hyp.sp <- function(d){
  hyp.size <- 2^d
  as.numeric(strsplit(substr(paste(as.integer(intToBits(2^d-1)), collapse=""),1,d
      ),"")[[1]])
  C <- sapply(0:(2^d-1), function(i) as.numeric(strsplit(substr(paste(as.integer(
      intToBits(i)), collapse=""),1,d),"")[[1]]))
  dim(C)
  dimnames(C) <- list(component=seq(d), hypothesis=seq(hyp.size))
```

```
  return(C)
}

#Generate a hypothesis with a certain number of errors out of the k 1-components
get.hyp.hit <- function(mu0, num.hit){
  d <- length(mu0)
  k <- sum(mu0)
  num.miss <- k-num.hit
  indx.1 <- which(mu0==1)
  indx.0 <- which(mu0==0)


  mu <- apply(num.miss, c("hypothesis","repetition"), function(miss){
      indx.one <- as.numeric(sample(as.character(indx.1), size=miss, replace=F))
      indx.zero <- as.numeric(sample(as.character(indx.0), size=miss, replace=F))
      res <- mu0
      res[indx.one] <- 0
      res[indx.zero] <- 1
      return(res)
  })
  names(dimnames(mu))[1] <- "component"
  dimnames(mu)[[1]] <- 1:d

  return(mu)
}

#Generate restricted hypothesis spaces
get.res.hyp.sp <- function(d, k){

  Cs <- lapply(k, function(ak) sapply(1:choose(d,ak), function(i) get.hyp.i(i, d,
      ak)))
  names(Cs) <- k
  return(Cs)
}

#Generate restricted hypothesis spaces
get.res.hyp.sp2 <- function(C, k){
  Cs <- lapply(k, function(num1) C[,which(apply(C,2,sum)==num1)])
  names(Cs) <- k
  return(Cs)
}

#Simlate Data
sim.data <- function(n, noise.level, dim.data, dimnames.data){
  d <- dim.data["component"]
  num.noise <- dim.data["noise"]
  reps <- dim.data["repetition"]
  num.spars <- dim.data["sparsity"]

  if("Zs" %in% names(dim.data)){
    num.Zs <- dim.data["Zs"]
  } else{
    num.Zs <- 1
  }

  #simulate normal data
  data.norm <- rnorm(num.spars*d*reps*num.Zs*num.noise,0,rep(noise.level/sqrt(n),
      rep(num.spars*d*reps*num.Zs, num.noise)))
  data.arr <- array(data.norm,dim=dim.data, dimnames=dimnames.data)
  return(data.arr)
}

#Create cost array
cost.mean.n2.old <- function(mu, mu0, X){

  a <- 1
  b <- 2
  res1 <- apply(X, c("repetition","Zs","noise.level"), function(vec) a*t(mu)%*%(1-
      b*vec-2*mu0))
  a <- 2
```

```
  b ← 1
  Y ← (X[,,1,]+X[,,2,])
  res2 ← apply(Y, c("repetition","noise.level"), function(vec) a*t(mu)%*%(1-b*vec
      -2*mu0))
  dimnames ← dimnames(res1)
  dimnames$Zs ← c(dimnames$Zs, "nZ12")
  res ← abind(res1[,1,],res1[,2,],res2, along=3)
  res ← aperm(res, c(1,3,2))
  dimnames(res) ← dimnames
  return(res)

}

cost.mean.n2 ← function(mu, mu0, X){

  res ← apply(X, c("repetition","Zs","noise.level"), function(vec) t(mu- (vec+mu0
      ))%*%(mu- (vec+mu0)))
    dimnames ← dimnames(res)
    dimnames$Zs ← c(dimnames$Zs, "nZ12")
  res ← abind(res[,1,],res[,2,],res[,1,]+res[,2,], along=3)
  res ← aperm(res, c(1,3,2))
  dimnames(res) ← dimnames
  return(res)
}

cost.mean.n1 ← function(mu, mu0, X){

  d ← length(mu0)
  res ← apply(X, c("repetition","Zs","noise.level"), function(vec) t(abs(mu- (vec
      +mu0)))%*%rep(1, d))
    dimnames ← dimnames(res)
    dimnames$Zs ← c(dimnames$Zs, "nZ12")
  res ← abind(res[,1,],res[,2,],res[,1,]+res[,2,], along=3)
  res ← aperm(res, c(1,3,2))
  dimnames(res) ← dimnames
  return(res)
}


#Create cost array sample case
# squared loss
cost.mean.smpl.n2 ← function(C, mu0, X){


  rep.Z ← length(dimnames(X)$repetition)
  num.noise ← length(dimnames(X)$noise.level)

  res ← sapply(1:rep.Z, function(i) apply(X[,i,,], c("Zs","noise.level"),
      function(vec) t(C[,i]- (vec+mu0))%*%(C[,i]- (vec+mu0))), simplify="array")
  res ← aperm(res, c(3,1,2))
  names(dimnames(res))[1] ← "repetition"
  dimnames(res)[[1]] ← 1:rep.Z


    dimnames ← dimnames(res)
    dimnames$Zs ← c(dimnames$Zs, "nZ12")
  res ← abind(res[,1,],res[,2,],res[,1,]+res[,2,], along=3)
  res ← aperm(res, c(1,3,2))
  dimnames(res) ← dimnames
  return(res)

}
# absolute loss
cost.mean.smpl.n1 ← function(C, mu0, X){


  rep.Z ← length(dimnames(X)$repetition)
  num.noise ← length(dimnames(X)$noise.level)
  d ← length(dimnames(X)$component)

```

```
  res ← sapply(1:rep.Z, function(i) apply(X[,i,,], c("Zs","noise.level"),
      function(vec) t(abs(C[,i]- (vec+mu0)))%*%rep(1,d)), simplify="array")
  res ← aperm(res, c(3,1,2))
  names(dimnames(res))[1] ← "repetition"
  dimnames(res)[[1]] ← 1:rep.Z


    dimnames ← dimnames(res)
    dimnames$Zs ← c(dimnames$Zs, "nZ12")
  res ← abind(res[,1,],res[,2,],res[,1,]+res[,2,], along=3)
  res ← aperm(res, c(1,3,2))
  dimnames(res) ← dimnames
  return(res)

}


#Create cost array importance sampling case
# squared loss
cost.mean.imp.smpl.n2 ← function(C, mu0, X){


  rep.Z ← length(dimnames(X)$repetition)
  num.noise ← length(dimnames(X)$noise.level)
  rep.mu ← length(dimnames(C)$hypothesis)
  num.Zs ← length(dimnames(X)$Zs)


  res ← sapply(1:rep.Z, function(i) sapply(1:num.Zs, function(z)
                              sapply(1:num.noise, function(n)
                                          sapply(1:rep.mu, function(hyp){
                                            t(C[,hyp,i]- (X[,i,z,n]+mu0))%*%(C[,
                                                hyp,i]- (X[,i,z,n]+mu0))
   }, simplify="array"), simplify="array"), simplify="array"),simplify="array")



  dimnames(res) ← list(hypothesis=1:rep.mu, noise.level=seq(num.noise), Zs=c("Z1"
      ,"Z2"), repetition=seq(rep.Z))
    dimnames ← dimnames(res)
    dimnames$Zs ← c(dimnames$Zs, "nZ12")
  res ← abind(res[,,1,],res[,,2,],res[,,1,]+res[,,2,], along=4)


  res ← aperm(res, c(1,2,4,3))
  dimnames(res) ← dimnames
  return(res)

}
# absolute loss
cost.mean.imp.smpl.n1 ← function(C, mu0, X){


  rep.Z ← length(dimnames(X)$repetition)
  num.noise ← length(dimnames(X)$noise.level)
  rep.mu ← length(dimnames(C)$hypothesis)
  num.Zs ← length(dimnames(X)$Zs)
  d ← length(dimnames(X)$component)


  res ← sapply(1:rep.Z, function(i) sapply(1:num.Zs, function(z)
                              sapply(1:num.noise, function(n)
                                          sapply(1:rep.mu, function(hyp){
                                            t(abs(C[,hyp,i]- (X[,i,z,n]+mu0)))%*%
                                                rep(1,d)
   }, simplify="array"), simplify="array"), simplify="array"),simplify="array")



```

```
236   dimnames(res) ← list(hypothesis=1:rep.mu, noise.level=seq(num.noise), Zs=c("Z1"
          ,"Z2"), repetition=seq(rep.Z))
237     dimnames ← dimnames(res)
238     dimnames$Zs ← c(dimnames$Zs, "nZ12")
239   res ← abind(res[,,1,],res[,,2,],res[,,1,]+res[,,2,], along=4)
240
241
242   res ← aperm(res, c(1,2,4,3))
243   dimnames(res) ← dimnames
244   return(res)
245
246 }
247
248
249 #Calculate boltzman weight, probabiity and partition arrays
250 # without log-sum-exp trick
251 get.arrays ← function(Cs, data.arr, mu0, beta, cost.fun){
252
253
254   d ← length(dimnames(data.arr)$component)
255   num.betas ← length(beta)
256   k ← as.numeric(dimnames(data.arr)$sparsity)
257   num.spars ← length(k)
258   hyp.size ← choose(d,k)
259   max.hyp.size ← 2^d
260
261
262
263   Rs ← lapply(1:num.spars, function(sp){
264       R ← sapply(as.data.frame(Cs[[sp]]),function(hyp) cost.fun(hyp,mu0[,sp],data
              .arr[sp,,,,]), simplify="array")
265       R ← aperm(R,c(4,1,2,3))
266       names(dimnames(R))[[1]] ← "hypothesis"
267       return(R)
268   })
269
270
271   #Calculate boltzman weights
272   ws ← lapply(Rs, function(el) {
273     w ← outer(el, beta, FUN=function(x,y) exp(-y*x))
274     names(dimnames(w))[length(names(dimnames(w)))] ← "beta"
275     dimnames(w)[[length(names(dimnames(w)))]] ← seq(num.betas)
276     return(w)
277   })
278
279   names(ws) ← 1:num.spars
280
281   #Calculate partition functions
282   Z ← sapply(ws, function(el) apply(el, c("repetition","Zs","noise.level","beta")
          ,sum), simplify="array")
283   Z ← aperm(Z, c(5,1,2,3,4))
284   names(dimnames(Z))[1] ← "sparsity"
285   dimnames(Z)[[1]] ← k
286
287
288   #Calculate gibbs probability distribution over hypothesis space
289   #prs ← lapply(1:num.spars, function(i){
290   #
291   # pr ← sapply(1:hyp.size[i], function(j) ws[[i]][j,,1:2,,]/Z[i,,1:2,,],
          simplify="array")
292   # pr ← aperm(pr, c(5,1,2,3,4))
293   # dimnames ← dimnames(ws[[i]])
294   # dimnames$Zs ← dimnames$Zs[1:2]
295   # dimnames(pr) ← dimnames
296   # return(pr)
297   #})
298
299   #names(prs) ← seq(num.spars)
300
301
```

```
302   res ← list(Rs=Rs, ws=ws, Z=Z) #, prs=prs
303 }
304 # with log-sum-exp trick
305 get.uf.arrays ← function(Cs, data.arr, mu0, beta, cost.fun){
306
307   #cost.fun ← cost.mean.n2
308   d ← length(dimnames(data.arr)$component)
309   num.betas ← length(beta)
310   k ← as.numeric(dimnames(data.arr)$sparsity)
311   num.spars ← length(k)
312   hyp.size ← choose(d,k)
313   max.hyp.size ← 2^d
314
315
316
317   Rs ← lapply(1:num.spars, function(sp){
318       R ← sapply(as.data.frame(Cs[[sp]]),function(hyp) cost.fun(hyp,mu0[,sp],data
              .arr[sp,,,,]), simplify="array")
319       R ← aperm(R,c(4,1,2,3))
320       names(dimnames(R))[[1]] ← "hypothesis"
321       return(R)
322   })
323
324
325
326   #Calculate -beta*R so that we may apply log-sum-exp trick to deal with
          underflow
327
328   #Calculate boltzman weights
329   as ← lapply(Rs, function(el) {
330     a ← outer(el, beta, FUN=function(x,y) -y*x)
331     names(dimnames(a))[length(names(dimnames(a)))] ← "beta"
332     dimnames(a)[[length(names(dimnames(a)))]] ← seq(num.betas)
333     return(a)
334   })
335
336   names(as) ← 1:num.spars
337
338
339   #Calculate boltzman weights
340   logZ ← sapply(as, function(el){
341     res ← apply(el, c("noise.level","Zs","repetition","beta"), function(vec){
342       b ← max(vec)
343       return(log(sum(exp(vec-b)))+b)
344       })
345     return(res)
346   }, simplify="array")
347
348   logZ ← aperm(logZ, c(5,1,2,3,4))
349   names(dimnames(logZ))[1] ← "sparsity"
350   dimnames(logZ)[[1]] ← k
351
352   res ← list(Rs=Rs, logZ=logZ) #, prs=prs
353 }
354
355
356 #Calculate boltzman weight, probabiity and partition arrays - sample version
357 # without log-sum-exp trick
358 get.smpl.arrays ← function(reps.mu, data.arr, mu0, beta, cost.fun){
359
360
361   d ← length(dimnames(data.arr)$component)
362   reps.Z ← length(dimnames(data.arr)$repetition)
363   num.betas ← length(beta)
364   k ← as.numeric(dimnames(data.arr)$sparsity)
365   num.spars ← length(k)
366   hyp.size ← choose(d,k)
367   max.hyp.size ← 2^d
368
369   # generate a list with the possible hypothesis numbers that can be sampled
```

```
370   hyp.no.list ← sapply(hyp.size, function(size) 1:size)
371
372
373   smpls ← sapply(1:num.spars, function(sp) sapply(1:reps.Z, function(i) sample(
          hyp.no.list[[sp]],size=reps.mu, replace=T) ), simplify="array")
374   dimnames(smpls) ← list(hypothesis=1:reps.mu, repetition=1:reps.Z, sparsity=1:
          num.spars)
375
376
377   R ← sapply(1:num.spars, function(sp){
378       R.aux ← sapply(1:reps.mu, function(j){
379         C ← sapply(smpls[j,,sp], function(i) get.hyp.i(i, d, k[sp]))
380         dimnames(C) ← list(component=1:d, repetition=1:reps.Z)
381         return(cost.fun(C,mu0[,sp],data.arr[sp,,,,]))
382       }, simplify="array")
383       R.aux ← aperm(R.aux,c(4,1,2,3))
384       names(dimnames(R.aux))[1] ← "hypothesis"
385       dimnames(R.aux)[[1]] ← 1:reps.mu
386       return(R.aux)
387   }, simplify="array")
388
389   R ← aperm(R, c(5,1,2,3,4))
390   names(dimnames(R))[1] ← "sparsity"
391   dimnames(R)[[1]] ← k
392
393
394
395   #Calculate boltzman weights
396   w ← outer(R, beta, FUN=function(x,y) exp(-y*x))
397   names(dimnames(w))[length(names(dimnames(w)))] ← "beta"
398   dimnames(w)[[length(names(dimnames(w)))]] ← seq(num.betas)
399
400   #Calculate partition functions
401   Z ← apply(w, c("sparsity","repetition","Zs","noise.level","beta"),mean)#*hyp.
          size
402
403   #Calculate gibbs probability distribution over hypothesis space
404   #pr ← sapply(1:reps.mu, function(i) w[,i,,1:2,,]/(Z[,,1:2,,]*reps.Z), simplify
          ="array")
405   #pr ← aperm(pr, c(1,6,2,3,4,5))
406   #dimnames ← dimnames(w)
407   #dimnames$Zs ← dimnames$Zs[1:2]
408   #dimnames(pr) ← dimnames
409
410   res ← list(R=R, w=w, Z=Z) #, pr=pr
411 }
412 # with log-sum-exp trick
413 get.smpl.uf.arrays ← function(reps.mu, data.arr, mu0, beta, cost.fun){
414
415   #cost.fun ← cost.mean.smpl.n2
416   d ← length(dimnames(data.arr)$component)
417   reps.Z ← length(dimnames(data.arr)$repetition)
418   num.betas ← length(beta)
419   k ← as.numeric(dimnames(data.arr)$sparsity)
420   num.spars ← length(k)
421   hyp.size ← choose(d,k)
422   max.hyp.size ← 2^d
423
424   # generate a list with the possible hypothesis numbers that can be sampled
425   hyp.no.list ← sapply(hyp.size, function(size) 1:size)
426
427
428   smpls ← sapply(1:num.spars, function(sp) sapply(1:reps.Z, function(i) sample(
          hyp.no.list[[sp]],size=reps.mu, replace=T) ), simplify="array")
429   dimnames(smpls) ← list(hypothesis=1:reps.mu, repetition=1:reps.Z, sparsity=1:
          num.spars)
430
431
432   R ← sapply(1:num.spars, function(sp){
433       R.aux ← sapply(1:reps.mu, function(j){
```

```
434         C ← sapply(smpls[j,,sp], function(i) get.hyp.i(i, d, k[sp]))
435         dimnames(C) ← list(component=1:d, repetition=1:reps.Z)
436         return(cost.fun(C,mu0[,sp],data.arr[sp,,,,]))
437       }, simplify="array")
438       R.aux ← aperm(R.aux,c(4,1,2,3))
439       names(dimnames(R.aux))[1] ← "hypothesis"
440       dimnames(R.aux)[[1]] ← 1:reps.mu
441       return(R.aux)
442   }, simplify="array")
443
444   R ← aperm(R, c(5,1,2,3,4))
445   names(dimnames(R))[1] ← "sparsity"
446   dimnames(R)[[1]] ← k
447
448   #Calculate boltzman weights
449   a ← outer(R, beta, FUN=function(x,y) -y*x)
450   names(dimnames(a))[length(names(dimnames(a)))] ← "beta"
451   dimnames(a)[[length(names(dimnames(a)))]] ← seq(num.betas)
452
453
454   #Calculate boltzman weights
455   logZ ← apply(a, c("sparsity","noise.level","Zs","repetition","beta"), function(
          vec){
456       b ← max(vec)
457       res ← log(sum(exp(vec-b)))+b
458       return(res)
459     })
460
461   logZ ← logZ - log(reps.mu)
462
463   res ← list(R=R, logZ=logZ)
464 }
465
466
467 #Calculate boltzman weight, probabiity and partition arrays - importance sampling
        version
468 # without log-sum-exp trick
469 get.imp.smpl.arrays ← function(reps.mu, data.arr, mu0, beta, cost.fun){
470
471   #cost.fun ← cost.mean.imp.smpl.n2
472   d ← length(dimnames(data.arr)$component)
473   reps.Z ← length(dimnames(data.arr)$repetition)
474   num.betas ← length(beta)
475   k ← as.numeric(dimnames(data.arr)$sparsity)
476   num.spars ← length(k)
477   hyp.size ← choose(d,k)
478   max.hyp.size ← 2^d
479
480   # generate a list with the possible hypothesis numbers that can be sampled
481   #hyp.no.list ← sapply(hyp.size, function(size) 1:size)
482
483   names(beta) ← 1:length(beta)
484
485   weights ← lapply(k, function(i){
486     hits.rng ← max(0,2*i-d):i
487     res ← rep(1/length(hits.rng), length(hits.rng))
488     names(res) ← hits.rng
489     return(res)
490   })
491
492   num.type ← lapply(k, function(i){
493     hits.rng ← max(0,2*i-d):i
494     res ← choose(i,hits.rng)*choose(d-i, i-(hits.rng))
495     names(res) ← hits.rng
496     return(res)
497   })
498
499
500   #lapply(weights, function(el) apply(el,"beta", sum))
501
```

```
smpls ← sapply(1:num.spars, function(sp)  sapply(1:reps.Z, function(i)
    as.numeric(sample(names(weights[[sp]]),size=reps.mu, replace=T, prob=
        weights[[sp]])), simplify="array"), simplify="array")

dimnames(smpls) ← list(hypothesis=1:reps.mu, repetition=1:reps.Z, sparsity=k)

#get the weights of each sampled hypothesis for when we calculate the weighted
    mean Z

reweights ← sapply(1:reps.mu, function(hyp)
                  sapply(1:reps.Z, function(i) sapply(1:num.spars,
                                  function(sp){
                                      indx ← match(smpls[hyp,i,sp],names(
                                          weights[[sp]]))
                                      (1/reps.mu)*(1/hyp.size[sp])*(num.type[[
                                          sp]][indx]/weights[[sp]][indx])
                                   }, simplify="array"), simplify="array"),
                                       simplify="array")

dimnames(reweights) ← list(sparsity=k, repetition=1:reps.Z, hypothesis=1:reps.
    mu)

reweights ← aperm(reweights, c(1,3,2))

C ← sapply(1:num.spars, function(sp){
   #print(sp)
   get.hyp.hit(mu0=mu0[,sp], num.hit=smpls[,,sp])
 }, simplify="array")
C ← aperm(C, c(4,1,2,3))
names(dimnames(C))[1] ← "sparsity"
dimnames(C)[[1]] ← k


#apply(C,c("sparsity","hypothesis","repetition"), function(vec) sum(vec))
#apply(C[2,,,],c("hypothesis","repetition"), function(vec) t(vec)%*%mu0[,2])==
    smpls[,,2]



R ← sapply(1:num.spars, function(sp) cost.fun(C=C[sp,,,], mu0=mu0[,sp], X=data.
    arr[sp,,,,]), simplify="array")



R ← aperm(R, c(5,1,2,3,4))
names(dimnames(R))[1] ← "sparsity"
dimnames(R)[[1]] ← k

#apply(R, c("sparsity","noise.level"), summary)

#Calculate boltzman weights
w ← outer(R, beta, FUN=function(x,y) exp(-y*x))
names(dimnames(w))[length(names(dimnames(w)))] ← "beta"
dimnames(w)[[length(names(dimnames(w)))]] ← seq(num.betas)



#Calculate partition functions
Z ← sapply(1:num.noise, function(noi)
              sapply(1:3, function(z)
                      sapply(1:num.betas, function(bet)
                                  w[,,noi,z,,bet]*reweights
                                    ,simplify="array"), simplify="array"),
                                        simplify="array")

names(dimnames(Z))[(length(names(dimnames(Z)))-2):length(names(dimnames(Z)))] ←
     c("beta","Zs","noise.level")
dimnames(Z)[[(length(names(dimnames(Z)))-2)]] ← seq(num.betas)
dimnames(Z)[[(length(names(dimnames(Z)))-1)]] ← c("Z1","Z2","nZ12")
dimnames(Z)[[(length(names(dimnames(Z))))]] ← seq(num.noise)
```

```
562
563   Z ← apply(Z, c("sparsity","repetition","Zs","noise.level","beta"),sum)#*hyp.
          size
564   Z ← aperm(Z, c(1,4,3,2,5))
565
566   #Calculate gibbs probability distribution over hypothesis space
567   pr ← sapply(1:reps.mu, function(i) w[,i,,1:2,,]/(Z[,,1:2,,]*reps.Z), simplify="
          array")
568   pr ← aperm(pr, c(1,6,2,3,4,5))
569   dimnames ← dimnames(w)
570   dimnames$Zs ← dimnames$Zs[1:2]
571   dimnames(pr) ← dimnames
572
573   res ← list(R=R, w=w, pr=pr, Z=Z)
574 }
575 # with log-sum-exp trick
576 get.imp.smpl.uf.arrays ← function(reps.mu, data.arr, mu0, beta, cost.fun){
577
578   #cost.fun ← cost.mean.imp.smpl.n2
579   d ← length(dimnames(data.arr)$component)
580   reps.Z ← length(dimnames(data.arr)$repetition)
581   num.betas ← length(beta)
582   k ← as.numeric(dimnames(data.arr)$sparsity)
583   num.spars ← length(k)
584   hyp.size ← choose(d,k)
585   max.hyp.size ← 2^d
586
587   # generate a list with the possible hypothesis numbers that can be sampled
588   #hyp.no.list ← sapply(hyp.size, function(size) 1:size)
589
590   names(beta) ← 1:length(beta)
591
592   weights ← lapply(k, function(i){
593     hits.rng ← max(0,2*i-d):i
594     res ← rep(1/length(hits.rng), length(hits.rng))
595     names(res) ← hits.rng
596     return(res)
597   })
598
599   num.type ← lapply(k, function(i){
600     hits.rng ← max(0,2*i-d):i
601     res ← choose(i,hits.rng)*choose(d-i, i-(hits.rng))
602     names(res) ← hits.rng
603     return(res)
604   })
605
606
607   #lapply(weights, function(el) apply(el,"beta", sum))
608
609   smpls ← sapply(1:num.spars, function(sp)  sapply(1:reps.Z, function(i)
610       as.numeric(sample(names(weights[[sp]]),size=reps.mu, replace=T, prob=
              weights[[sp]])), simplify="array"), simplify="array")
611
612   dimnames(smpls) ← list(hypothesis=1:reps.mu, repetition=1:reps.Z, sparsity=k)
613
614   #get the weights of each sampled hypothesis for when we calculate the weighted
          mean Z
615
616   reweights ← sapply(1:reps.mu, function(hyp)
617                      sapply(1:reps.Z, function(i) sapply(1:num.spars,
618                                         function(sp){
619                                             indx ← match(smpls[hyp,i,sp],names(
                                                    weights[[sp]]))
620                                             return((1/reps.mu)*(1/hyp.size[sp])*(num.
                                                    type[[sp]][indx]/weights[[sp]][indx])
                                                    )
621                                          }, simplify="array"), simplify="array"),
                                               simplify="array")
622
```

```
623   dimnames(reweights) ← list(sparsity=k, repetition=1:reps.Z, hypothesis=1:reps.
          mu)
624
625   reweights ← aperm(reweights, c(1,3,2))
626
627   #lets check that the reweights add up to one, or 1/reps.mu in this case since
          we already added that
628 # apply(reweights, c("sparsity", "repetition"), sum); 1/reps.mu
629
630
631   C ← sapply(1:num.spars, function(sp){
632      #print(sp)
633      get.hyp.hit(mu0=mu0[,sp], num.hit=smpls[,,sp])
634    }, simplify="array")
635   C ← aperm(C, c(4,1,2,3))
636   names(dimnames(C))[1] ← "sparsity"
637   dimnames(C)[[1]] ← k
638
639
640   #apply(C,c("sparsity","hypothesis","repetition"), function(vec) sum(vec))
641   #apply(C[2,,,],c("hypothesis","repetition"), function(vec) t(vec)%*%mu0[,2])==
          smpls[,,2]
642
643
644
645   R ← sapply(1:num.spars, function(sp) cost.fun(C=C[sp,,,], mu0=mu0[,sp], X=data.
          arr[sp,,,,]), simplify="array")
646
647   R ← aperm(R, c(5,1,2,3,4))
648   names(dimnames(R))[1] ← "sparsity"
649   dimnames(R)[[1]] ← k
650   #apply(R, c("sparsity","noise.level"), summary)
651
652
653
654   #Calculate -beta*R so that we may apply log-sum-exp trick to deal with
          underflow
655
656   a ← outer(R, beta, FUN=function(x,y) -y*x)
657   names(dimnames(a))[length(names(dimnames(a)))] ← "beta"
658   dimnames(a)[[length(names(dimnames(a)))]] ← seq(num.betas)
659   b ← apply(a, c("sparsity","noise.level","Zs","repetition","beta"),max)
660
661
662   #Calculate boltzman weights
663   w ← apply(a, c("sparsity","noise.level","Zs","repetition","beta"), function(vec
          ) exp(vec-max(vec)))
664   w ← aperm(w, c(2,1,3,4,5,6))
665
666   #Calculate partition functions
667   Z ← sapply(1:num.noise, function(noi)
668                 sapply(1:3, function(z)
669                         sapply(1:num.betas, function(bet)
670                                     w[,,noi,z,,bet]*reweights
671                                       ,simplify="array"), simplify="array"),
                                           simplify="array")
672
673   names(dimnames(Z))[(length(names(dimnames(Z)))-2):length(names(dimnames(Z)))] ←
          c("beta","Zs","noise.level")
674   dimnames(Z)[[(length(names(dimnames(Z)))-2)]] ← seq(num.betas)
675   dimnames(Z)[[(length(names(dimnames(Z)))-1)]] ← c("Z1","Z2","nZ12")
676   dimnames(Z)[[(length(names(dimnames(Z))))]] ← seq(num.noise)
677
678   Z ← apply(Z, c("sparsity","repetition","Zs","noise.level","beta"),sum)#*hyp.
          size
679   Z ← aperm(Z, c(1,4,3,2,5))
680
681   logZ ← log(Z) + b
682
683
```

```
  res ← list(R=R, w=w, Z=Z, logZ=logZ)
}

#calculate information content, average info. content and generalization capacity
# exhaustive algo.
get.info ← function(hyp.size, logZ){

  cardinality ←  array(hyp.size, dim=dim(logZ[,,1,,]), dimnames=dimnames(logZ
      [,,1,,]))

  num.noise ← length(dimnames(logZ)$noise.level)
  #Calculate Information Content
  info.content ← log(cardinality) + logZ[,,3,,] - logZ[,,1,,] -logZ[,,2,,]

  #Calculate average information content
  avg.info.content ← apply(info.content, c("sparsity","noise.level","beta"), mean
      )

  #Calculate standard deviation information content
  sd.info.content ← apply(info.content, c("sparsity","noise.level","beta"), sd)

  #Obtain generalization capacity and corresponding beta
  indx.beta.opt ← apply(avg.info.content,c("sparsity","noise.level"),which.max)
  gen.capacity ← apply(avg.info.content,c("sparsity","noise.level"),max)

  indx.arr ← melt(indx.beta.opt)
  names(indx.arr)[3] ← "beta"
  indx.arr$sparsity ← match(indx.arr$sparsity, k)
  indx.arr$sd ← sd.info.content[as.matrix(indx.arr)]
  sd.gen.capacity ← cast(indx.arr, sparsity~noise.level, value="sd")
  sd.gen.capacity ← as.matrix(sd.gen.capacity[,2:dim(sd.gen.capacity)[2]])
  dimnames(sd.gen.capacity) ← dimnames(gen.capacity)

  return(list(info.content=info.content, avg.info.content=avg.info.content, sd.
      info.content=sd.info.content, indx.beta.opt=indx.beta.opt ,gen.capacity=gen
      .capacity, sd.gen.capacity=sd.gen.capacity))
}

#calculate information content, average info. content and generalization capacity
# sampling algo.
get.smpl.info ← function(hyp.size, logZ){

  k ← as.numeric(dimnames(logZ)$sparsity)
  num.noise ← length(dimnames(logZ)$noise.level)
  #Calculate Information Content
  info.content ←   logZ[,,3,,] - logZ[,,1,,] -logZ[,,2,,]

  #Calculate average information content
  avg.info.content ← apply(info.content, c("sparsity","noise.level","beta"), mean
      )

  #Calculate standard deviation information content
  sd.info.content ← apply(info.content, c("sparsity","noise.level","beta"), sd)

  #Obtain generalization capacity and corresponding beta
  indx.beta.opt ← apply(avg.info.content,c("sparsity","noise.level"),which.max)
  gen.capacity ← apply(avg.info.content,c("sparsity","noise.level"),max)

  indx.arr ← melt(indx.beta.opt)
  names(indx.arr)[3] ← "beta"
  indx.arr$sparsity ← match(indx.arr$sparsity, k)
  indx.arr$sd ← sd.info.content[as.matrix(indx.arr)]
  sd.gen.capacity ← cast(indx.arr, sparsity~noise.level, value="sd")
  sd.gen.capacity ← as.matrix(sd.gen.capacity[,2:dim(sd.gen.capacity)[2]])
  dimnames(sd.gen.capacity) ← dimnames(gen.capacity)

  return(list(info.content=info.content, avg.info.content=avg.info.content, sd.
      info.content=sd.info.content, indx.beta.opt=indx.beta.opt ,gen.capacity=gen
      .capacity, sd.gen.capacity=sd.gen.capacity))
}
```

```
747
748 # importance sampling algo.
749 get.imp.smpl.info ← function(hyp.size, logZ){
750
751   k ← as.numeric(dimnames(logZ)$sparsity)
752   num.noise ← length(dimnames(logZ)$noise.level)
753   #Calculate Information Content
754   info.content ←   logZ[,,3,,] - logZ[,,1,,] -logZ[,,2,,] #+ log(hyp.size)
755
756   #Calculate average information content
757   avg.info.content ← apply(info.content, c("sparsity","noise.level","beta"), mean
          )
758
759   #Calculate standard deviation information content
760   sd.info.content ← apply(info.content, c("sparsity","noise.level","beta"), sd)
761
762   #Obtain generalization capacity and corresponding beta
763   indx.beta.opt ← apply(avg.info.content,c("sparsity","noise.level"),which.max)
764   gen.capacity ← apply(avg.info.content,c("sparsity","noise.level"),max)
765
766   indx.arr ← melt(indx.beta.opt)
767   names(indx.arr)[3] ← "beta"
768   indx.arr$sparsity ← match(indx.arr$sparsity, k)
769   indx.arr$sd ← sd.info.content[as.matrix(indx.arr)]
770   sd.gen.capacity ← cast(indx.arr, sparsity~noise.level, value="sd")
771   sd.gen.capacity ← as.matrix(sd.gen.capacity[,2:dim(sd.gen.capacity)[2]])
772   dimnames(sd.gen.capacity) ← dimnames(gen.capacity)
773
774   return(list(info.content=info.content, avg.info.content=avg.info.content, sd.
          info.content=sd.info.content, indx.beta.opt=indx.beta.opt ,gen.capacity=gen
          .capacity, sd.gen.capacity=sd.gen.capacity))
775 }
776
777 #plot avg info. content vs beta for different noise levels
778 plot.avg.info.content ← function(info, beta, noise.level, choose.noise, choose.
        sparsity){
779
780   choose.num.noise ← length(choose.noise)
781   choose.num.sparsity ← length(choose.sparsity)
782   k ← as.numeric(dimnames(info$info.content)$sparsity)
783
784
785
786   for(j in 1:choose.num.sparsity){
787   plot(beta, info$avg.info.content[choose.sparsity[j],choose.noise[1],], type="l"
          , ylim=c(-0.5,max(info$avg.info.content[choose.sparsity[j],choose.noise,])*
          1.1), xlab="beta",ylab="Average Information Content", main=paste("sparsity
          = ",k[choose.sparsity[j]]))
788     for(i in 1:choose.num.noise){
789       lines(beta, info$avg.info.content[choose.sparsity[j],choose.noise[i],], col
              =i)
790       lines(x=beta[info$indx.beta.opt[choose.sparsity[j],choose.noise[i]]], info$
              gen.capacity[choose.sparsity[j],choose.noise[i]], type="p", col=i)
791       text(x=beta[info$indx.beta.opt[choose.sparsity[j],choose.noise[i]]],y=info$
              gen.capacity[choose.sparsity[j], choose.noise[i]]+0.1,labels=round(
              noise.level[choose.noise[i]],1),cex=0.8)
792     }
793   }
794   abline(h=0, col="grey")
795
796   for(i in 1:choose.num.noise){
797   plot(beta, info$avg.info.content[choose.sparsity[1],choose.noise[i],], type="l"
          , ylim=c(-0.5,max(info$avg.info.content[choose.sparsity,choose.noise[i],])*
          1.1), xlab="beta",ylab="Average Information Content", main=paste("noise
          level = ",round(noise.level[choose.noise[i]],1)))
798   for(j in 1:choose.num.sparsity){
799       lines(beta, info$avg.info.content[choose.sparsity[j],choose.noise[i],], col
              =j)
800       lines(x=beta[info$indx.beta.opt[choose.sparsity[j],choose.noise[i]]], info$
              gen.capacity[choose.sparsity[j],choose.noise[i]], type="p", col=j)
```

```
      text(x=beta[info$indx.beta.opt[choose.sparsity[j],choose.noise[i]]],y=info$
          gen.capacity[choose.sparsity[j], choose.noise[i]]+0.1,labels=k[choose.
          sparsity[j]],cex=0.8)
    }
  }
  abline(h=0,col="grey")

}

#plot generalization capacity vs noise level
plot.gen.capacity <- function(n, noise.level, prs, info, Cs, mu0, choose.sparsity)
    {

  d <- dim(Cs[[1]])[1]

  reps <- length(dimnames(prs[[1]])$repetition)
  num.noise <- length(noise.level)
  bayes.hit   <- (pnorm(sqrt(n)/(2*noise.level)))^d
  bayes.error <- 1-bayes.hit
  num.spars <- length(prs)

  indx.mat <- melt(info$indx.beta.opt)
  colnames(indx.mat)[3] <- "beta"
  pr.opt <- lapply(1:num.spars, function(i) apply(prs[[i]], c("hypothesis","
      repetition","Zs"), function(mat) mat[as.matrix(indx.mat[which(indx.mat$
      sparsity==i),c("noise.level","beta")])]))
  names(pr.opt) <- 1:num.spars
  pr.opt <- lapply(pr.opt, function(el){
     names(dimnames(el))[1] <- "noise.level"
     dimnames(el)[[1]] <- 1:num.noise
     el <- aperm(el, c(2,3,4,1))
     return(el)
   })
   indx.true.hyp <- sapply(1:num.spars, function(i) which(apply(Cs[[i]]==mu0[,i],
       "hypothesis", all)))
   pr.gibbs.true <- sapply(1:num.spars, function(i) apply(pr.opt[[i]][indx.true.
       hyp[i],,1,], c("noise.level"), mean))
  names(dimnames(pr.gibbs.true)) <- c("noise.level","sparsity")
  dimnames(pr.gibbs.true)$noise.level <- 1:num.noise
  dimnames(pr.gibbs.true)$sparsity <- 1:num.spars



  plot(noise.level, pr.opt[[choose.sparsity]][indx.true.hyp[choose.sparsity
      ],1,1,], type="l",ylim=range(pr.opt[[choose.sparsity]][indx.true.hyp[choose
      .sparsity],,,]))
  for(i in 1:reps) for(j in 1:2) lines(noise.level, pr.opt[[choose.sparsity]][
      indx.true.hyp[choose.sparsity],i,j,])
  lines(noise.level, pr.gibbs.true[,choose.sparsity], col="blue")

  y.gen <- pretty(c(0,info$gen.capacity[choose.sparsity,]))
   # add extra room to the left of the plot
   par(oma=c(0,2,0,0))
   #generalization capacity
   plot(x=noise.level, y=info$gen.capacity[choose.sparsity,], col='blue', type='l
       ', ylim=range(y.gen), main=paste("Generalization capacity by noise level,
       sparsity= ", k[choose.sparsity]), xlab='noise level', ylab='',xaxt='n',
       yaxt='n', lwd=0.75)
   abline(h=0, col="grey")
   axis(2, col='blue', at=y.gen, labels=y.gen)
   #gibbs probability
   par(new=T)
   plot(x=noise.level, y=pr.gibbs.true[,choose.sparsity], ylim=c(0,1),col='red',
       type='l', lwd=0.75,xaxt='n', axes=F, ylab='')
   axis(side=2, at=y.gen/max(y.gen), labels=round(y.gen/max(y.gen),1), col='red',
        line=2)
   #1-bayes error
   par(new=T)
   plot(x=noise.level, y=bayes.hit, ylim=c(0,1),col='green', type='l', lwd=0.75,
       xaxt='n', axes=F, ylab='')
```

```
855     axis(side=2, at=y.gen/max(y.gen), labels=rep("",length(y.gen)), col='green',
            line=2.1)
856     #legend
857     legend(x=0, y=.8, legend=c('gen. cap', 'sd. info cont','gibbs true', '1-bayes
            err.'), col=c("blue","black","red","green"), lwd=c(3.5,  3.5, 3.5))
858     #standar deviation information content
859     par(new=T)
860     plot(x=noise.level, y=info$sd.gen.capacity[choose.sparsity,], ylim=c(0,max(
            info$sd.gen.capacity[choose.sparsity,])),col='black', type='l', lwd=0.75,
            xaxt='n', axes=F, ylab='')
861     axis(side=2, at=y.gen/max(y.gen)*max(info$sd.gen.capacity[choose.sparsity,]),
            labels=round(y.gen/max(y.gen)*max(info$sd.gen.capacity)), col='black',
            line=4)
862     #x-axis
863     axis(side=1, at=noise.level,   labels=round(noise.level,1))
864     abline(v=noise.level, col='grey', lwd=0.5)
865
866     #All sparsities
867      y.gen <- pretty(c(0,gen.capacity))
868      # add extra room to the left of the plot
869      par(oma=c(0,2,0,0))
870      #generalization capacity
871      plot(x=noise.level, y=info$gen.capacity[1, ], col='blue', type='l', ylim=
             range(y.gen), main="Generalization capacity by noise level and sparsity",
              xlab='noise level', ylab='', lwd=0.75)
872     for(i in 1:num.spars) lines(x=noise.level, y=info$gen.capacity[i, ], col=i)
873     #x-axis
874      axis(side=1, at=noise.level,   labels=round(noise.level,1))
875      abline(v=noise.level, col='grey', lwd=0.5)
876
877      legend(x=10, y=5, legend=paste("sparsity = ", 1:num.spars), col=1:num.spars,
            lwd=c(3.5,  3.5, 3.5))
878
879
880 }
881
882 #plot generalization capacity vs noise level
883 plot.smpl.gen.capacity <- function(d, n, noise.level, info){
884
885
886   reps.Z <- length(dimnames(info$info.content)$repetition)
887   num.noise <- length(noise.level)
888   bayes.hit   <- (pnorm(sqrt(n)/(2*noise.level)))^d
889   bayes.error <- 1-bayes.hit
890
891
892   y.gen <- pretty(c(0,info$gen.capacity))
893    # add extra room to the left of the plot
894    par(oma=c(0,2,0,0))
895    #generalization capacity
896    plot(x=noise.level, y=info$gen.capacity, col='blue', type='l', ylim=range(y.
           gen), main="Generalization capacity by noise level", xlab='noise level',
           ylab='',xaxt='n', yaxt='n', lwd=0.75)
897    #segments(x0=noise.level, y0=info$gen.capacity-info$sd.gen.capacity, x1 =
           noise.level, y1 = info$gen.capacity+info$sd.gen.capacity, col="grey", lwd
           =3)
898
899    abline(h=0, col="grey")
900    axis(2, col='blue', at=y.gen, labels=y.gen)
901    #1-bayes error
902    par(new=T)
903    plot(x=noise.level, y=bayes.hit, ylim=c(0,1),col='green', type='l', lwd=0.75,
           xaxt='n', axes=F, ylab='')
904    axis(side=2, at=y.gen/max(y.gen), labels=rep("",length(y.gen)), col='green',
           line=2.1)
905
906    legend(x=0, y=.8, legend=c('gen. cap', 'sd. info cont', '1-bayes err.'), col=c
           ("blue","black","green"), lwd=c(3.5,  3.5, 3.5))
907    #standar deviation information content
908    par(new=T)
```

```
    plot(x=noise.level, y=info$sd.gen.capacity, ylim=c(0,max(info$sd.gen.capacity)),col='black', type='l', lwd=0.75,xaxt='n', axes=F, ylab='')
    axis(side=2, at=y.gen/max(y.gen)*max(info$sd.gen.capacity), labels=round(y.gen/max(y.gen)*max(info$sd.gen.capacity)), col='black', line=4)

    #x-axis
    axis(side=1, at=noise.level,   labels=round(noise.level,1))
    abline(v=noise.level, col='grey', lwd=0.5)

}

#obtain generalization capacity and standard deviation of information content
get.gen.capacity <- function(data.arr, Cs, mu0, beta, cost.fun){


  hyp.size <- sapply(Cs, function(x) dim(x)[2])


  arrays <- get.uf.arrays(Cs, data.arr, mu0, beta, cost.fun)
  info <- get.info(hyp.size, arrays$logZ)
  return(list(arrays=arrays, info=info))
}

#obtain generalization capacity and standard deviation of information content - sample version
get.smpl.gen.capacity <- function(reps.mu, data.arr, mu0, beta, cost.fun){


  d <- dim(data.arr)["component"]
  k <- as.numeric(dimnames(data.arr)$sparsity)
  hyp.size <- choose(d,k)

  arrays <- get.smpl.uf.arrays(reps.mu, data.arr, mu0, beta, cost.fun)
  info <- get.smpl.info(hyp.size, logZ=arrays$logZ)
  return(list(arrays=arrays, info=info))
}

#obtain generalization capacity and standard deviation of information content - importance sample version
get.imp.smpl.gen.capacity <- function(reps.mu, data.arr, mu0, beta, cost.fun){


  d <- dim(data.arr)["component"]
  k <- as.numeric(dimnames(data.arr)$sparsity)
  hyp.size <- choose(d,k)
  arrays <- get.imp.smpl.uf.arrays(reps.mu, data.arr, mu0, beta, cost.fun)
  info <- get.imp.smpl.info(hyp.size, logZ=arrays$logZ)
  return(list(arrays=arrays, info=info))
}
```

A.2 Script

```
library(ggplot2)
library(reshape)
library(abind)

source("./functions.R")

################################################################################
# Set Parameters
################################################################################

n <- 100
reps.Z <- 100
reps.mu <- 1000
d <- 10
```

```
k ← c(1,4)
num.spars ← length(k)
hyp.size ← choose(d,k)
max.hyp.size ← 2^d
set.seed(1)
C ← get.hyp.sp(d)
Cs ← get.res.hyp.sp(d, k)
mu0 ← sapply(Cs, function(mat) mat[,sample(1:ncol(mat),size=1)])            # the "true" localization parameter
dimnames(mu0) ← list(component=1:d, sparsity=k)

num.noise ← 20
noise.level ← seq(0.1,15,length.out=num.noise)
dim.data ← c(sparsity=num.spars,component=d,repetition=reps.Z,Zs=2, noise=num.noise)
dimnames.data ← list(sparsity=k, component=seq(d),repetition=seq(reps.Z),Zs=c("Z1","Z2"), noise.level=seq(num.noise))
num.betas ← 100
beta ← seq(1e-12, 100, length.out=num.betas)


################################################################################
# Script
################################################################################



#simulate data
set.seed(4)
data.arr ← sim.data(n, noise.level, dim.data, dimnames.data)

#calculate gen capacity

res.mean.n2          ← get.gen.capacity(data.arr, Cs, mu0, beta, cost.mean.n2)
res.smpl.mean.n2     ← get.smpl.gen.capacity(reps.mu, data.arr, mu0, beta, cost.mean.smpl.n2)
res.imp.smpl.mean.n2 ← get.imp.smpl.gen.capacity(reps.mu, data.arr, mu0, beta, cost.mean.imp.smpl.n2)

res.mean.n1          ← get.gen.capacity(data.arr, Cs, mu0, beta, cost.mean.n1)
res.smpl.mean.n1     ← get.smpl.gen.capacity(reps.mu, data.arr, mu0, beta, cost.mean.smpl.n1)
res.imp.smpl.mean.n1 ← get.imp.smpl.gen.capacity(reps.mu, data.arr, mu0, beta, cost.mean.imp.smpl.n1)

##########################################################
#Compare sampling and exhaustive methods for 1 d
##########################################################

################################
# R(mu,X)=sum(mu-X_bar)^2
################################

###################
# Gen. capacity
###################

choose.sparsity ← 2

max.gen ← max(c(res.mean.n2$info$gen.capacity[choose.sparsity,],res.smpl.mean.n2$info$gen.capacity[choose.sparsity,],res.imp.smpl.mean.n2$info$gen.capacity[choose.sparsity,]))
plot(noise.level,res.mean.n2$info$gen.capacity[choose.sparsity,], ylim=c(0,max.gen),col="black",pch=1,type="b", xlab="noise level", ylab="generalization capacity", main="generalization capacity vs. noise")
lines(noise.level,res.smpl.mean.n2$info$gen.capacity[choose.sparsity,], col="blue", type="b",pch=2)
lines(noise.level,res.imp.smpl.mean.n2$info$gen.capacity[choose.sparsity,], col="green", type="b",pch=3)
```

```
legend('topright', legend=c('exhaustive','sampling','imp. sampling'), col=c("
    black","blue","green"),pch=c(1,2,3),lwd=c(3.5,  3.5, 3.5))


##################
# SD Info Content
##################

max.gen <- max(c(res.mean.n2$info$sd.gen.capacity[choose.sparsity,],res.smpl.mean.
    n2$info$sd.gen.capacity[choose.sparsity,],res.imp.smpl.mean.n2$info$sd.gen.
    capacity[choose.sparsity,]))
plot(noise.level,res.mean.n2$info$sd.gen.capacity[choose.sparsity,], ylim=c(0,max
    .gen),col="black",pch=1,type="b", xlab="noise level", ylab="std dev info",
    main="std. dev. info. content vs. noise")
lines(noise.level,res.smpl.mean.n2$info$sd.gen.capacity[choose.sparsity,], col="
    blue", type="b",pch=2)
lines(noise.level,res.imp.smpl.mean.n2$info$sd.gen.capacity[choose.sparsity,],
    col="green", type="b",pch=3)
legend('topright', legend=c('exhaustive','sampling','imp. sampling'), col=c("
    black","blue","green"),pch=c(1,2,3),lwd=c(3.5,  3.5, 3.5))

################################
# R(mu,X)=sum(|mu-X_bar|)
################################


##################
# Gen. capacity
##################

choose.sparsity <- 2

max.gen <- max(c(res.mean.n1$info$gen.capacity[choose.sparsity,],res.smpl.mean.n1$
    info$gen.capacity[choose.sparsity,],res.imp.smpl.mean.n1$info$gen.capacity[
    choose.sparsity,]))
plot(noise.level,res.mean.n1$info$gen.capacity[choose.sparsity,], ylim=c(0,max.
    gen),col="black",pch=1,type="b", xlab="noise level", ylab="generalization
    capacity", main="generalization capacity vs. noise")
lines(noise.level,res.smpl.mean.n1$info$gen.capacity[choose.sparsity,], col="blue
    ", type="b",pch=2)
lines(noise.level,res.imp.smpl.mean.n1$info$gen.capacity[choose.sparsity,], col="
    green", type="b",pch=3)
legend('topright', legend=c('exhaustive','sampling','imp. sampling'), col=c("
    black","blue","green"),pch=c(1,2,3),lwd=c(3.5,  3.5, 3.5))


##################
# SD Info Content
##################

max.gen <- max(c(res.mean.n1$info$sd.gen.capacity[choose.sparsity,],res.smpl.mean.
    n1$info$sd.gen.capacity[choose.sparsity,],res.imp.smpl.mean.n1$info$sd.gen.
    capacity[choose.sparsity,]))
plot(noise.level,res.mean.n1$info$sd.gen.capacity[choose.sparsity,], ylim=c(0,max
    .gen),col="black",pch=1,type="b", xlab="noise level", ylab="std dev info",
    main="std. dev. info. content vs. noise")
lines(noise.level,res.smpl.mean.n1$info$sd.gen.capacity[choose.sparsity,], col="
    blue", type="b",pch=2)
lines(noise.level,res.imp.smpl.mean.n1$info$sd.gen.capacity[choose.sparsity,],
    col="green", type="b",pch=3)
legend('topright', legend=c('exhaustive','sampling','imp. sampling'), col=c("
    black","blue","green"),pch=c(1,2,3),lwd=c(3.5,  3.5, 3.5))

#######################################################
# Compare cost functions for exhaustive method only
#######################################################

choose.sparsity <- 2

```

```
max.gen ← max(c(res.mean.n1$info$gen.capacity[choose.sparsity,],res.mean.n2$info$
    gen.capacity[choose.sparsity,]))
plot(noise.level,res.mean.n2$info$gen.capacity[choose.sparsity,], ylim=c(0,max.
    gen),col="black",pch=1,type="b", xlab="noise level", ylab="generalization
    capacity", main="generalization capacity vs. noise")
lines(noise.level,res.mean.n1$info$gen.capacity[choose.sparsity,], col="blue",
    type="b",pch=2)
legend('topright', legend=c('squared loss','absolute loss'), col=c("black","blue"
    ),pch=c(1,2),lwd=c(3.5,  3.5))


##########################################################
#Compare sampling and exhaustive methods for a few d
##########################################################

#Now we obtain gen capacity and sd.gen capacity, for exhaustive and sampling
    methods and for noise = 1,2 and d=3,...,d.max

#Parameters
n ← 100
reps.Z ← 100
reps.mu ← 500
k ← c(1,4)
num.spars ← length(k)
noise.level ← c(2,4)
num.noise ← length(noise.level)
choose.sparsity ← 2
d.max ← 20
finish ← 0
ds ← 6:d.max
num.ds ← length(ds)
dimnames ← list(components=ds, noise=noise.level, sim.method=c("exhaustive","
    sampling","imp.samp"), cost.function=c("mean.n2","mean.n1"))
gen.capacity ← array(NA, dim=c(length(ds), num.noise,3,2) , dimnames=dimnames)
sd.gen.capacity ← array(NA,dim=c(length(ds), num.noise,3,2), dimnames=dimnames)


num.betas ← 100
beta.min ← 1e-12
beta.max ← 100
beta ← exp(seq(log(beta.min), log(beta.max), length.out=num.betas))
beta.opt.n1 ← matrix(NA, num.ds, num.noise)
beta.opt.n2 ← matrix(NA, num.ds, num.noise)


for(i in 1:length(ds)){
  start ← proc.time()[1]

  hyp.size ← choose(ds[i],k)
  max.hyp.size ← 2^d.max
  #set.seed(1)
  C ← get.hyp.sp(ds[i])
  Cs ← get.res.hyp.sp(ds[i], k)
  mu0 ← sapply(Cs, function(mat) mat[,sample(1:ncol(mat),size=1)])            #
      the "true" localization parameter
  dimnames(mu0) ← list(component=1:ds[i], sparsity=k)
  dim.data ← c(sparsity=num.spars,component=ds[i],repetition=reps.Z,Zs=2, noise=
      num.noise)
  dimnames.data ← list(sparsity=k, component=seq(ds[i]),repetition=seq(reps.Z),Zs
      =c("Z1","Z2"), noise.level=seq(num.noise))
  #set.seed(4)
  data.arr ← sim.data(n, noise.level, dim.data, dimnames.data)
  res.mean.n2          ← get.gen.capacity(data.arr, Cs, mu0, beta, cost.mean.n2)
  res.smpl.mean.n2     ← get.smpl.gen.capacity(reps.mu, data.arr, mu0, beta, cost
      .mean.smpl.n2)
  res.imp.smpl.mean.n2 ← get.imp.smpl.gen.capacity(reps.mu, data.arr, mu0, beta,
      cost.mean.imp.smpl.n2)
  res.mean.n1          ← get.gen.capacity(data.arr, Cs, mu0, beta, cost.mean.n1)
  res.smpl.mean.n1     ← get.smpl.gen.capacity(reps.mu, data.arr, mu0, beta, cost
      .mean.smpl.n1)
```

```
  res.imp.smpl.mean.n1 ← get.imp.smpl.gen.capacity(reps.mu, data.arr, mu0, beta,
      cost.mean.imp.smpl.n1)

  beta.opt.n2[i,] ← beta[res.imp.smpl.mean.n2$info$indx.beta.opt[choose.sparsity
      ,]]
  beta.opt.n1[i,] ← beta[res.imp.smpl.mean.n1$info$indx.beta.opt[choose.sparsity
      ,]]
  print("beta opt n2")
  print(beta.opt.n2[i,])
  print("beta opt n1")
  print(beta.opt.n1[i,])
  beta.min.aux ← min(beta.opt.n2[i,],beta.opt.n1[i,])*0.5
  beta.max.aux ← max(beta.opt.n2[i,],beta.opt.n1[i,])*1.5
  beta ← exp(seq(log(beta.min.aux), log(beta.max.aux), length.out=num.betas))
  print("new beta range")
  print(paste(round(beta.min.aux,4),"-",round(beta.max.aux,4), sep=""))

  gen.capacity[i,,1,1]    ← res.mean.n2$info$gen.capacity[choose.sparsity,]
  gen.capacity[i,,2,1]    ← res.smpl.mean.n2$info$gen.capacity[choose.sparsity,]
  gen.capacity[i,,3,1]    ← res.imp.smpl.mean.n2$info$gen.capacity[choose.
      sparsity,]
  sd.gen.capacity[i,,1,1] ← res.mean.n2$info$sd.gen.capacity[choose.sparsity,]
  sd.gen.capacity[i,,2,1] ← res.smpl.mean.n2$info$sd.gen.capacity[choose.sparsity
      ,]
  sd.gen.capacity[i,,3,1] ← res.imp.smpl.mean.n2$info$sd.gen.capacity[choose.
      sparsity,]

  gen.capacity[i,,1,2]    ← res.mean.n1$info$gen.capacity[choose.sparsity,]
  gen.capacity[i,,2,2]    ← res.smpl.mean.n1$info$gen.capacity[choose.sparsity,]
  gen.capacity[i,,3,2]    ← res.imp.smpl.mean.n1$info$gen.capacity[choose.
      sparsity,]
  sd.gen.capacity[i,,1,2] ← res.mean.n1$info$sd.gen.capacity[choose.sparsity,]
  sd.gen.capacity[i,,2,2] ← res.smpl.mean.n1$info$sd.gen.capacity[choose.sparsity
      ,]
  sd.gen.capacity[i,,3,2] ← res.imp.smpl.mean.n1$info$sd.gen.capacity[choose.
      sparsity,]

  finish ← c(finish, (proc.time()[1] - start))
  print("d")
  print(ds[i])
  print("finish")
  print(finish)
}

finish ← finish[2:length(finish)]
plot(ds, finish)

################################
# R(mu,X)=sum(mu-X_bar)^2
################################

###################
# Gen. capacity
###################

max.gen ← max(gen.capacity[,,,1])
plot(ds,gen.capacity[,1,1,1], ylim=c(0,max.gen),col="blue",pch=1,type="b", xlab="
    Dimension d", ylab="generalization capacity", main="generalization capacity
    vs. dimension")
lines(ds,gen.capacity[,2,1,1], col="blue", type="b",pch=2)
lines(ds,gen.capacity[,1,2,1], col="red", type="b",pch=1)
lines(ds,gen.capacity[,2,2,1], col="red", type="b",pch=2)
lines(ds,gen.capacity[,1,3,1], col="green", type="b",pch=1)
lines(ds,gen.capacity[,2,3,1], col="green", type="b",pch=2)
legend('topleft', legend=c(paste('noise  ',noise.level[1],'- exhaustive', sep="")
    ,paste('noise ',noise.level[2],' - exhaustive', sep=""),paste('noise ',noise.
    level[1],' - sampling', sep=""),paste('noise ',noise.level[2],' - sampling',
    sep=""),paste('noise ',noise.level[1],' - imp sampling', sep=""),paste('noise
     ',noise.level[2],' - imp sampling', sep="")), col=c("blue","blue","red","red
    ","green","green"),pch=c(1,2,1,2,1,2),lwd=rep(3.5,6))
```

```
229
230 ##################
231 # SD Info Content
232 ##################
233
234 max.gen ← max(sd.gen.capacity[,,,1])
235 plot(ds,sd.gen.capacity[,1,1,1], ylim=c(0,max.gen),col="blue",pch=1,type="b",
        xlab="Dimension d", ylab="std dev info", main="std. dev. info. content vs.
        dimension")
236 lines(ds,sd.gen.capacity[,2,1,1], col="blue", type="b",pch=2)
237 lines(ds,sd.gen.capacity[,1,2,1], col="red", type="b",pch=1)
238 lines(ds,sd.gen.capacity[,2,2,1], col="red", type="b",pch=2)
239 lines(ds,sd.gen.capacity[,1,3,1], col="green", type="b",pch=1)
240 lines(ds,sd.gen.capacity[,2,3,1], col="green", type="b",pch=2)
241 legend('topleft', legend=c(paste('noise  ',noise.level[1],'- exhaustive', sep="")
        ,paste('noise ',noise.level[2],' - exhaustive', sep=""),paste('noise ',noise.
        level[1],' - sampling', sep=""),paste('noise ',noise.level[2],' - sampling',
        sep=""),paste('noise ',noise.level[1],' - imp sampling', sep=""),paste('noise
         ',noise.level[2],' - imp sampling', sep="")), col=c("blue","blue","red","red
        ","green","green"),pch=c(1,2,1,2,1,2),lwd=rep(3.5,6))
242
243 ################################
244 # R(mu,X)=sum(|mu-X_bar|)
245 ################################
246
247 ##################
248 # Gen. capacity
249 ##################
250
251 max.gen ← max(gen.capacity[,,,2])
252 plot(ds,gen.capacity[,1,1,2], ylim=c(0,max.gen),col="blue",pch=1,type="b", xlab="
        Dimension d", ylab="generalization capacity", main="generalization capacity
        vs. dimension")
253 lines(ds,gen.capacity[,2,1,2], col="blue", type="b",pch=2)
254 lines(ds,gen.capacity[,1,2,2], col="red", type="b",pch=1)
255 lines(ds,gen.capacity[,2,2,2], col="red", type="b",pch=2)
256 lines(ds,gen.capacity[,1,3,2], col="green", type="b",pch=1)
257 lines(ds,gen.capacity[,2,3,2], col="green", type="b",pch=2)
258 legend('topleft', legend=c(paste('noise  ',noise.level[1],'- exhaustive', sep="")
        ,paste('noise ',noise.level[2],' - exhaustive', sep=""),paste('noise ',noise.
        level[1],' - sampling', sep=""),paste('noise ',noise.level[2],' - sampling',
        sep=""),paste('noise ',noise.level[1],' - imp sampling', sep=""),paste('noise
         ',noise.level[2],' - imp sampling', sep="")), col=c("blue","blue","red","red
        ","green","green"),pch=c(1,2,1,2,1,2),lwd=rep(3.5,6))
259
260 ##################
261 # SD Info Content
262 ##################
263
264 max.gen ← max(sd.gen.capacity[,,,2])
265 plot(ds,sd.gen.capacity[,1,1,2], ylim=c(0,max.gen),col="blue",pch=1,type="b",
        xlab="Dimension d", ylab="std dev info", main="std. dev. info. content vs.
        dimension")
266 lines(ds,sd.gen.capacity[,2,1,2], col="blue", type="b",pch=2)
267 lines(ds,sd.gen.capacity[,1,2,2], col="red", type="b",pch=1)
268 lines(ds,sd.gen.capacity[,2,2,2], col="red", type="b",pch=2)
269 lines(ds,sd.gen.capacity[,1,3,2], col="green", type="b",pch=1)
270 lines(ds,sd.gen.capacity[,2,3,2], col="green", type="b",pch=2)
271 legend('topleft', legend=c(paste('noise  ',noise.level[1],'- exhaustive', sep="")
        ,paste('noise ',noise.level[2],' - exhaustive', sep=""),paste('noise ',noise.
        level[1],' - sampling', sep=""),paste('noise ',noise.level[2],' - sampling',
        sep=""),paste('noise ',noise.level[1],' - imp sampling', sep=""),paste('noise
         ',noise.level[2],' - imp sampling', sep="")), col=c("blue","blue","red","red
        ","green","green"),pch=c(1,2,1,2,1,2),lwd=rep(3.5,6))
272
273 ######################################################
274 # Compare cost functions for exhaustive method only
275 ######################################################
276
277 max.gen ← max(gen.capacity[,,1,])
```

```
278 plot(ds,gen.capacity[,1,1,1], ylim=c(0,max.gen),col="black",pch=1,type="b", xlab=
        "Dimension d", ylab="generalization capacity", main="generalization capacity
        vs. dimension")
279 lines(ds,gen.capacity[,1,1,2], col="blue", type="b",pch=2)
280 lines(ds,gen.capacity[,2,1,1], col="black", type="b",pch=2)
281 lines(ds,gen.capacity[,2,1,2], col="blue", type="b",pch=2)
282 legend('topleft', legend=c(paste('noise  ',noise.level[1],'- squared loss', sep="
        "),paste('noise ',noise.level[2],' - squared loss', sep=""),paste('noise ',
        noise.level[1],' - absolute loss', sep=""),paste('noise ',noise.level[2],' -
        absolute loss', sep="")), col=c("black","black","blue","blue"),pch=c(1,2,1,2)
        ,lwd=rep(3.5,4))
283
284
285
286 ##############################################################################
287 # Now we obtain gen capacity and sd.gen capacity, for sampling methods and for
288 # noise = 1,2 and d=3,...,10000
289 ##############################################################################
290
291 # Now we obtain gen capacity and sd.gen capacity, for exhaustive and sampling
292 # methods and for noise = 1,2 and d=3,...,d.max
293
294 #Parameters
295 n ← 100
296 reps.Z ← 100
297 reps.mu ← 100
298 k ← c(1,4)
299 num.spars ← length(k)
300 noise.level ← rep(8,2)
301 num.noise ← length(noise.level)
302 choose.sparsity ← 2
303 d.max ← 10000
304 finish ← 0
305 num.ds ← 100
306 ds ← round(seq(10,d.max, length.out=num.ds))
307 dimnames ← list(components=ds, noise=noise.level, cost.function=c("mean.n2","mean
        .n1"))
308 gen.capacity ← array(NA, dim=c(length(ds), num.noise, 2) , dimnames=dimnames)
309 sd.gen.capacity ← array(NA,dim=c(length(ds), num.noise, 2), dimnames=dimnames)
310
311
312 num.betas ← 100
313 beta.min ← 1e-12
314 beta.max ← 50
315 beta ← exp(seq(log(beta.min), log(beta.max), length.out=num.betas))
316 beta.opt.n1 ← matrix(NA, num.ds, num.noise)
317 beta.opt.n2 ← matrix(NA, num.ds, num.noise)
318
319 for(i in 1:length(ds)){
320   start ← proc.time()[1]
321
322   mu0 ← sapply(k, function(ak){
323      res ← rep(0,ds[i])
324      indx ← sample(1:ds[i],size=ak)
325      res[indx] ← 1
326      return(res)
327    })
328
329   dimnames(mu0) ← list(component=1:ds[i], sparsity=k)
330   dim.data ← c(sparsity=num.spars,component=ds[i],repetition=reps.Z,Zs=2, noise=
          num.noise)
331   dimnames.data ← list(sparsity=k, component=seq(ds[i]),repetition=seq(reps.Z),Zs
          =c("Z1","Z2"), noise.level=seq(num.noise))
332   #set.seed(4)
333   data.arr ← sim.data(n, noise.level, dim.data, dimnames.data)
334
335   res.imp.smpl.mean.n2 ← get.imp.smpl.gen.capacity(reps.mu, data.arr, mu0, beta,
          cost.mean.imp.smpl.n2)
336   res.imp.smpl.mean.n1 ← get.imp.smpl.gen.capacity(reps.mu, data.arr, mu0, beta,
          cost.mean.imp.smpl.n1)
```

```
337
338   beta.opt.n2[i,] ← beta[res.imp.smpl.mean.n2$info$indx.beta.opt[choose.sparsity
          ,]]
339   beta.opt.n1[i,] ← beta[res.imp.smpl.mean.n1$info$indx.beta.opt[choose.sparsity
          ,]]
340   print("beta opt n2")
341   print(beta.opt.n2[i,])
342   print("beta opt n1")
343   print(beta.opt.n1[i,])
344   beta.min.aux ← min(beta.opt.n2[i,],beta.opt.n1[i,])*0.5
345   beta.max.aux ← max(beta.opt.n2[i,],beta.opt.n1[i,])*1.5
346   beta ← exp(seq(log(beta.min.aux), log(beta.max.aux), length.out=num.betas))
347   print("new beta range")
348   print(paste(round(beta.min.aux,4),"-",round(beta.max.aux,4), sep=""))
349
350
351   gen.capacity[i,,1]    ← res.imp.smpl.mean.n2$info$gen.capacity[choose.sparsity
          ,]
352   sd.gen.capacity[i,,1] ← res.imp.smpl.mean.n2$info$sd.gen.capacity[choose.
          sparsity,]
353   gen.capacity[i,,2]    ← res.imp.smpl.mean.n1$info$gen.capacity[choose.sparsity
          ,]
354   sd.gen.capacity[i,,2] ← res.imp.smpl.mean.n1$info$sd.gen.capacity[choose.
          sparsity,]
355
356   finish ← c(finish, (proc.time()[1] - start))
357
358   print("gen capacity")
359   print(gen.capacity[i,,1])
360
361   print("d")
362   print(ds[i])
363   print("finish")
364   print(finish[length(finish)])
365 }
366
367 finish ← finish[2:length(finish)]
368 plot(ds, finish)
369
370 plot(rep(ds,2), as.numeric(beta.opt.n2),ylim=range(c(as.numeric(beta.opt.n2),as.
        numeric(beta.opt.n1))))
371 plot(rep(ds,2), as.numeric(beta.opt.n1),ylim=range(c(as.numeric(beta.opt.n2),as.
        numeric(beta.opt.n1))))
372
373 ################################
374 # R(mu,X)=sum(mu-X_bar)^2
375 ################################
376
377 ###################
378 # Gen. capacity
379 ###################
380
381 max.gen ← max(gen.capacity[,,1])
382 x ← rep(ds, num.noise)
383 y ← as.numeric(gen.capacity[,,1])
384 indx ← order(x)
385 x ← x[indx]
386 gc.mean.n2 ← y[indx]
387 plot(x,gc.mean.n2, ylim=c(0,max.gen),col="blue",pch=1,type="p", xlab="Dimension d
        ", ylab="generalization capacity", main="generalization capacity vs.
        dimension")
388 lines(ds,log(choose(ds,k[choose.sparsity])), col="red", type="l")
389 gc.spl.mean.n2 ← smooth.spline(x,gc.mean.n2,penalty=4)
390 lines(gc.spl.mean.n2, col="black")
391 legend('bottomright', legend=c(paste('noise ', noise.level[1], ' - imp sampling',
         sep=""),"trend",paste('log(choose(d,',k[choose.sparsity], '))',sep="")), col
        =c("blue","black","red"),pch=c(1,0,0),lwd=rep(3.5,3))
392
393 ###################
394 # SD Info Content
```

```
##################

max.gen <- max(sd.gen.capacity[,,1])
x <- rep(ds, num.noise)
y <- as.numeric(sd.gen.capacity[,,1])
indx <- order(x)
x <- x[indx]
sdgc.mean.n2 <- y[indx]
plot(x,sdgc.mean.n2, ylim=c(0,max.gen),col="blue",pch=1,type="p", xlab="Dimension d", ylab="sd generalization capacity", main="std. dev. info. content vs. dimension")
sdgc.spl.mean.n2 <- smooth.spline(x,sdgc.mean.n2,penalty=4)
lines(sdgc.spl.mean.n2, col="black")
legend('bottomright', legend=c(paste('noise ', noise.level[1], ' - imp sampling', sep=""),"trend"), col=c("blue","black"),pch=c(1,0),lwd=rep(3.5,2))

###############################
# R(mu,X)=sum(|mu-X_bar|)
###############################

##################
# Gen. capacity
##################

max.gen <- max(gen.capacity[,,2])
x <- rep(ds, num.noise)
y <- as.numeric(gen.capacity[,,2])
indx <- order(x)
x <- x[indx]
gc.mean.n1 <- y[indx]
plot(x,gc.mean.n1, ylim=c(0,max.gen),col="blue",pch=1,type="p", xlab="Dimension d", ylab="generalization capacity", main="generalization capacity vs. dimension")
lines(ds,log(choose(ds,k[choose.sparsity])), col="red", type="l")
gc.spl.mean.n1 <- smooth.spline(x,gc.mean.n1,penalty=4)
lines(gc.spl.mean.n1, col="black")
legend('bottomright', legend=c(paste('noise ', noise.level[1], ' - imp sampling', sep=""),"trend",paste('log(choose(d,',k[choose.sparsity], '))',sep="")), col=c("blue","black","red"),pch=c(1,0,0),lwd=rep(3.5,3))

##################
# SD Info Content
##################

max.gen <- max(sd.gen.capacity[,,2])
x <- rep(ds, num.noise)
y <- as.numeric(sd.gen.capacity[,,2])
indx <- order(x)
x <- x[indx]
sdgc.mean.n1 <- y[indx]
plot(x,sdgc.mean.n1, ylim=c(0,max.gen),col="blue",pch=1,type="p", xlab="Dimension d", ylab="sd generalization capacity", main="std. dev. info. content vs. dimension")
sdgc.spl.mean.n1 <- smooth.spline(x,sdgc.mean.n1,penalty=4)
lines(sdgc.spl.mean.n1, col="black")
legend('bottomright', legend=c(paste('noise ', noise.level[1], ' - imp sampling', sep=""),"trend"), col=c("blue","black"),pch=c(1,0),lwd=rep(3.5,2))

#######################################################
# Compare cost functions for exhaustive method only
#######################################################

max.gen <- max(gen.capacity)
plot(x,gc.mean.n2, ylim=c(0,max.gen),col="black",pch=1,type="p", xlab="Dimension d", ylab="generalization capacity", main="generalization capacity vs. dimension")
lines(x,gc.mean.n1,col="blue",pch=2,type="p")
lines(gc.spl.mean.n2, col="black",lwd=3)
lines(gc.spl.mean.n1, col="blue",lwd=3)
```

```
452 legend('bottomright', legend=c(paste('noise ', noise.level[1], '  squared loss',
        sep=""),paste('noise ', noise.level[1], '  absolute loss', sep="")), col=c("
        black","blue"),pch=c(1,2),lwd=rep(3.5,2))
```